Ph.D. Program in Information and Communication Technology

Curriculum Emerging computing models: algorithms, software architectures and intelligent systems

XXXVIII cycle

Advanced optimization models and methods for electrical mobility

SSD: MATH-06/A

Doctoral Program Supervisor
Prof. Davide Di Ruscio

Candidate
Isabella Presutti Gasbarro

Advisor
Prof. Claudio Arbib

Co-Advisors
Dott. Andrea Manno
Prof. Mustafa Pinar
Dott. Andrea Pizzuti

A.Y. 2024/2025

C., N. & M.

Acknowledgements

I would first like to acknowledge my main supervisor, prof. Claudio Arbib, for accompanying me throughout the PhD program and for providing the institutional framework within which this research was carried out.

I am sincerely grateful to my co-supervisors. In particular, I would like to thank Andrea Manno for his support in the machine learning aspects of this work and for his help during moments in which the research process became challenging, and Andrea Pizzuti for his invaluable guidance and supervision, especially during the period spent at Bilkent University, which represented a particularly meaningful phase of my doctoral journey.

I gratefully acknowledge Atlante srl for co-funding my PhD program and, in particular, Paolo Testa, for his support, trust, and for the many lessons that have contributed significantly to my academic growth.

I would like to thank my friends, both within and outside the PhD program, who were also colleagues with whom I shared meaningful moments, discussions, and experiences. Their presence and willingness to listen were essential, especially during the most difficult periods.

I am deeply grateful to my parents for their constant support and for the trust they have always placed in me, particularly on an emotional level. Their encouragement has been fundamental throughout this journey.

Finally, I would like to thank my fiancé, Claudio. Without the need for many words, he is the reason why I am here today writing these acknowledgements with confidence in myself, in my abilities, and with my mental well-being still intact. Forever.

Isabella Presutti Gasbarro, L'Aquila, December 2025

Abstract

Electric vehicle (EV) adoption is rapidly increasing, yet the development of charging infrastructure still faces strategic and operational challenges. Identifying profitable locations for new charging stations and designing effective pricing policies are central issues for both private operators and public planners. Existing works formalize the problem employing different methodologies, e.g., flow capturing model with capacity constraints, or pricing policy decisions. Nevertheless, to the best of our knowledge, none of them combine the aforementioned approaches with multi-period pricing and heterogeneous user preferences at both spatial and price levels. In addition, it is crucial to incorporate user behaviour derived from real charging session data rather than relying solely on declared preferences from surveys. Therefore, the contribution of this dissertation is two-fold. As the first contribution, we constructed a binary classification model to predict whether candidate locations are likely to exhibit high daily charging activity and determine which features have a major impact on station usage. The analysis of EV charging demand was carried out using a real-world dataset integrating spatial, temporal, and infrastructural attributes. Results show that the predictive model achieves an accuracy of 70% in identifying high-demand areas, providing an interpretable and computationally efficient tool for preliminary siting decisions. As the second and main contribution, leveraging this analysis, we propose a bilevel capacitated charging station location and pricing model. In the upper level, a charging point operator determines where to build new stations, how many connectors to install, and how to set multi-period prices. In the lower level, users select stations by minimizing an individual utility function that reflects price sensitivity, distance costs, and location-specific attractiveness. The attractiveness component is informed by the empirical insights derived from the aforementioned machine learning analysis, thereby capturing demand response and the impact of pricing on customer station choice. Computational experiments were conducted both on randomly generated and realistic instances that demonstrated the viability of the approach and highlighted the impact of temporal demand distribution and charging duration on instance hardness. Overall, the combination of the machine learning techniques and optimization components provides a comprehensive decision-support framework for EV infrastructure planning, enabling operators to (i) anticipate demand patterns and (ii) design profitable, user-oriented charging networks.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Electric vehicles (EVs) are increasingly regarded as a viable and promising alternative to conventional internal combustion engine (ICE) vehicles, offering substantial advantages at both the individual and societal levels. From an individual perspective, EVs provide lower operating costs, primarily due to reduced fuel expenditures and less frequent maintenance interventions, as electric powertrains contain fewer moving parts and are subject to reduced mechanical wear. They also offer the potential for longer service life and an enhanced driving experience, characterized by quieter operation, smoother acceleration, and the immediate availability of torque. In addition, the progressive diversification of EV models across market segments has broadened consumer appeal, making electric mobility accessible to a wider range of users.

From a collective standpoint, the widespread diffusion of electric vehicles plays a pivotal role in addressing multiple environmental and societal challenges. EV adoption contributes to the reduction of local air pollutants—such as nitrogen oxides (NO_x) and particulate matter (PM)—which are typically associated with ICE vehicles and are responsible for respiratory illnesses and declining urban air quality. Moreover, considering that the transportation sector is one of the major contributors to global CO_2 emissions, the electrification of mobility is imperative for long-term environmental sustainability and for achieving the climate neutrality objectives set by international agreements. The shift towards EVs therefore represents not only a technological transformation but also a necessary societal transition towards cleaner, more resilient mobility systems.

However, the sustained growth of the EV market does not depend solely on vehicle availability, technological performance, or consumer awareness. A crucial factor influencing EV adoption is the strategic development of a reliable, dense, and user-friendly charging infrastructure. Range anxiety, long charging times, and uncertainty about charger availability remain among the most frequently cited barriers to EV uptake. According to the EAFO Consumer Monitor 2023, 28% of surveyed drivers report a perceived shortage of public slow chargers, while an additional 18% indicate that the number of fast chargers remains insufficient [3]. These findings highlight the persistent gap between user expectations and the actual charging network deployment, reinforcing the need for improved planning and targeted investments.

A viable strategy to tackle these challenges involves simultaneously optimizing (i) the spatial distribution of EV charging stations, (ii) their technological configuration (e.g., number of connectors and power levels), and (iii) the associated pricing schemes offered to end-users. Unlike many other service facilities—such as retail stores or leisure venues—charging and refueling stations are often visited as intermediate stops rather than as final destinations. Drivers are typically willing to make short deviations along their routes to access charging facilities, making location planning a problem that must account for both origin–destination mobility patterns and deviation-based user behaviors. Furthermore, competition among Charge Point Operators (CPOs) and the increasing heterogeneity of charging technologies (AC, DC, and ultra-fast systems) motivate the integration of pricing policies into the planning process. An effective pricing strategy must account for users' purchasing power, their sensitivity to waiting times and connector power, and potential reactions to competitors' offerings.

In this context, optimization models—particularly those integrating routing decisions, capacity considerations, and real-world behavioral aspects—play a central role in supporting the efficient design of EV charging networks. Such models enable planners to identify cost-effective solutions that balance infrastructure investment, operational constraints, and user satisfaction, ultimately contributing to the accelerated adoption of electric mobility and to the broader sustainability goals of modern transportation systems.

1.1 State of the Art on Optimization Models for Charging Station Location

The problem of optimally locating electric vehicle (EV) charging infrastructure has been extensively studied across operations research, transportation science, and energy systems engineering. As EV adoption accelerates, methodological contributions have evolved to incorporate increasingly realistic representations of user behaviour, traffic flows, energy constraints, and market interactions. The literature now features a wide spectrum of mathematical models that differ in their assumptions, objectives, and structural complexity, yet share the common goal of supporting strategic planning decisions for charging infrastructure deployment.

Broadly speaking, existing approaches can be classified into several major families, including flow-based models, facility location models with refueling constraints, deviation-based formulations, capacity-aware models, and bilevel or equilibrium-based frameworks. Among these, *flow-based models* constitute a foundational class, as they explicitly account for origin–destination (OD) traffic patterns and thus capture the spatial distribution of mobility demand. These formulations have served as the basis for many subsequent extensions that incorporate range limitations, stochasticity, user equilibrium, or economic considerations.

In what follows, we first review the class of *flow-capturing models*, which represent the earliest and most fundamental line of research in this domain. We then discuss their evolution into *flow-refueling location models*, which introduce vehicle autonomy constraints and ensure trip feasibility. Finally, we examine *deviation-based models*, which enhance behavioural realism by allowing drivers to deviate from their shortest paths in order to access charging services. This progression highlights the conceptual and methodological development of location models for EV charging stations, moving from simple coverage-based logic toward models capable of capturing the complexities of real-world mobility and energy demand.

1.1.1 Flow-Capturing Models

Flow-capturing models (FCM) represent one of the earliest and most influential approaches to locating refueling or charging stations in transportation networks. These models are built on the premise that the demand for charging arises along origin–destination (OD) traffic flows, and that infrastructure should therefore be placed at locations that *intercept* as much of this flow as possible. Unlike traditional demand-based facility location problems—which assume that users travel from demand points directly to the nearest facility—flow-capturing models recognize that charging events often occur as intermediate stops rather than final destinations.

In their classical form, flow-capturing models assume that drivers follow predetermined shortest paths and that a vehicle is "captured" if its path intersects at least one charging station. This formulation allows planners to identify optimal facility placements that maximize the total intercepted flow or meet a required coverage threshold.

A classical application of the flow-capturing approach to EV charging infrastructure is provided by Cruz-Zambrano et al. [15], who apply the FCM to the city of Barcelona to determine strategic locations for fast charging stations. Their model relies on real OD traffic data and emphasizes the ability of flow-capturing formulations to handle dense urban networks where demand is highly directional and sensitive to mobility patterns. The authors show that capturing a sufficient proportion of daily flows can substantially improve accessibility while minimizing the number of required stations.

Building upon the same modeling philosophy, Tan and Lin [51] introduce stochasticity into the flow-capturing framework. Their stochastic flow-capturing location and allocation model accounts for uncertainty in EV traffic volumes and user behavior, enabling a more robust selection of charging locations under demand fluctuations. By integrating allocation decisions, the authors move beyond pure coverage and address how the captured flow should be distributed among selected facilities.

Wu and Sioshansi [55] further extend the stochastic flow-capturing paradigm by explicitly modeling uncertain EV flows and evaluating the impact of demand variability on network performance. Their model emphasizes that uncertainty is not only a source of risk but also a structural feature of evolving EV markets, where early-stage adoption patterns can significantly affect planning decisions.

Another notable development is the integration of wireless charging technologies. Riemann et al. [44] present a flow-capturing location model under stochastic user equilibrium conditions to determine optimal placements of wireless charging facilities along road networks. Their work is particularly relevant for long-distance mobility, where the ability to recharge dynamically along a route alters both the coverage logic and users' routing choices. The incorporation of user equilibrium also introduces behavioral realism, as drivers' path choices respond to travel time and charging availability.

Beyond single-level formulations, hierarchical planning frameworks have been proposed. Kong et al. [29] develop a hierarchical optimization model that coordinates the location of charging stations across multiple decision layers, considering both regional accessibility and local operational conditions. Their approach demonstrates the flexibility of flow-capturing principles when embedded into more elaborate planning structures.

Finally, the work of Honma and Kuby [23] provides a comparative perspective by examining node-based and path-based formulations for hydrogen refueling stations. Although their study is motivated by hydrogen infrastructure rather than EV charging, the underlying methodological insights apply directly. The authors show that path-based models (such as FCM and its variants) generally provide superior coverage of mobile demand, while node-based models may be more convenient but fail to fully capture the complexities of route-based behavior.

1.1.2 Flow-Refueling Location Models

While flow-capturing models focus on intercepting origin–destination (OD) flows along their shortest paths, they do not explicitly account for the limited driving range of vehicles. This limitation is critical in the context of alternative-fuel and electric vehicles, for which insufficient refueling or charging opportunities may prevent the completion of long-distance trips. To address this issue, the *Flow-Refueling Location Model* (FRLM) was introduced as a range-aware extension of flow-based location models. The FRLM seeks to determine the locations of a fixed number of refueling or charging stations such that the volume of OD flow that can be feasibly refueled, given vehicle range constraints, is maximized.

In their seminal work, Kuby and Lim [30] introduce the Flow-Refueling Location Model (FRLM) to account for vehicle range constraints when locating alternative-fuel stations. In this formulation, an OD flow is *covered* only if the installed stations along its path allow the vehicle to complete the trip without exceeding its driving range between successive refueling opportunities. This requirement yields a significantly more complex mixed-integer program than classical flow-capturing models, but provides a far more realistic representation of long-distance travel for limited-range vehicles.

Several extensions have adapted the FRLM to practical planning contexts. Upchurch et al. [53] incorporate capacity limits to create a *capacitated* FRLM, showing how congestion and limited service capability alter optimal deployment patterns. The authors also propose a revised objective based on vehicle-miles traveled, offering a more meaningful measure of covered demand.

Given the high computational burden of the FRLM, Lim and Kuby [35] develop heuristic algorithms—including greedy constructions and metaheuristics—that provide near-optimal solutions on realistically sized networks. Their work demonstrates that exploiting problem-specific structures can drastically reduce computation times without substantial loss of accuracy.

To further improve tractability, MirHassani and Ebrazi [38] propose a reformulation based on an expanded network representation. By embedding range feasibility directly into the network, their approach avoids enumerating all feasible refueling patterns, resulting in more compact models and improved scalability.

The FRLM has also been extended to multi-period settings. Chung and Kwon [13] formulate a multi-stage deployment model for EV charging stations on Korean expressways, combining FRLM-style range constraints with temporal investment decisions. Their results highlight the importance of planning infrastructure roll-out over time rather than relying solely on myopic strategies.

Decomposition techniques have been employed to handle additional realism and larger instances. Arslan and Karaşan [4] use Benders decomposition to address an FRLM-type problem for plug-in hybrids, separating location and flow feasibility decisions to achieve scalability. Similarly, Lee and Han [34] develop a Benders-and-Price method for probabilistic travel ranges, capturing uncertainty in effective driving distances.

Finally, Zhang et al. [61] embed FRLM logic within a coupled transportation–power network model to jointly optimize siting and sizing of fast-charging stations. Their integrated approach highlights the interdependence between EV mobility and power system constraints, extending the FRLM paradigm beyond pure transportation modeling.

1.1.3 Deviation Flow-Capturing Models

Although flow-refueling location models (FRLM) provide a realistic representation of range-feasible trips, they typically assume that drivers follow predetermined shortest paths. In practice, however, drivers may be willing to deviate from their preferred route to access a charging or refueling station, particularly when the deviation is small and the infrastructure network is sparse. To capture this behavioural realism, the class of *deviation-flow* models was developed. These models explicitly consider that OD flows may be partially or fully recaptured if facilities are located on alternative paths whose detour cost remains acceptable to users. As such, deviation-flow models bridge classical flow-based coverage and behaviourally informed location modelling, providing a more flexible representation of station accessibility under real-world conditions.

A pioneering formulation incorporating route deviation was introduced by Kim and Kuby [26], who proposed the *Deviation Flow Refueling Location Model* (DFRLM). In this model, drivers are allowed to choose alternative paths that require additional travel distance, and the captured demand decreases as the detour increases. Their approach relies on a pre-generated set of feasible deviation paths and introduces a decay function linking detour distance to the fraction of flow willing to deviate. This contribution established the conceptual foundation for modelling the trade-off between accessibility and detour tolerance.

Following this seminal work, Jiang et al. [24] introduced a *capacitated deviation-flow fueling location model*, extending the DFRLM to include station capacity constraints. Their model allocates flow to stations under both detour acceptance and service limitations, allowing for a more realistic representation of congestion effects at stations. This work constitutes one of the earliest attempts to jointly model deviation behaviour and capacity restrictions in the context of charging station planning.

To address the computational burden associated with generating large sets of deviation paths, several methodological innovations were proposed. Yildiz et al. [57] developed a path-segment based formulation that circumvents explicit path enumeration, thereby improving scalability on large networks. Their approach decomposes paths into elementary segments and constructs deviation-feasible paths implicitly within the model. Kim and Kuby [27] later proposed a network transformation heuristic for the DFRLM, transforming the original network into an augmented representation in which deviation effects can be captured through standard location algorithms. This heuristic offers a computationally efficient alternative to exact formulations.

Kinay et al. [31] further contributed by developing a full-cover deviation model in which drivers select optimal deviation routes that minimize the total en-route recharging need. Their model shifts the focus from pure path feasibility to optimal path selection conditioned on charging requirements, incorporating more detailed characteristics of EV energy consumption. Xu et al. [56] introduced a battery charging action-based path model that integrates nonlinear demand elasticity into the deviation process. Their approach accounts for user response to detour distance and charging decisions, thereby enhancing behavioural fidelity.

A second line of research incorporates station capacity and uncertainty more explicitly. Miralinaghi et al. [37] formulated a multi-period mixed-integer deviation-flow model that assumes constant travel times but introduces staircase operational costs reflecting economies of scale at stations. Their work highlights the importance of temporal dynamics and cost structures in the long-term planning of charging infrastructure. Guo et al. [22] proposed a leader–follower (Stackelberg-type) deviation-flow model that captures flow decay under distance deviations, enabling sensitivity analyses on behavioural parameters. Building on these ideas, Kinay et al. [32] and Yildiz et al. [58] examined capacitated deviation-flow problems under demand uncertainty, demonstrating how robust and stochastic modelling approaches can mitigate the risks associated with variable user behaviour and uncertain adoption patterns.

More recently, deviation-flow logic has been integrated into models that couple transportation networks with energy distribution systems. Sa'adati et al. [45] developed a joint allocation model for renewable energy sources (RESs) and plug-in electric vehicle (PEV) fast-charging stations on coupled transportation and distribution networks. Their formulation incorporates deviation-flow behaviour into both the siting of charging stations and the allocation of local renewable generation, illustrating how deviation-aware accessibility considerations must be balanced with operational constraints of the electricity grid.

1.1.4 Pricing Policies in EV Charging Infrastructure Planning

While the majority of classical location models emphasize spatial accessibility, range feasibility, and capacity management, the strategic role of pricing in EV charging infrastructure planning has received increasing attention in recent years. Pricing policies directly influence users' charging choices, their willingness to deviate from preferred routes, and their adoption of different connector power levels. For an operator, pricing determines revenue streams, affects station attractiveness relative to competing alternatives, and shapes long-term profitability. At the system level, pricing can also be used to balance demand across time and space, mitigate congestion, and encourage efficient use of installed capacity.

As the EV charging ecosystem becomes more heterogeneous—featuring multiple charge point operators, diverse pricing schemes (flat rates, dynamic pricing, energy-based tariffs, time-based tariffs), and varying power levels—integrating pricing decisions into optimization models has become crucial. The interplay between location, capacity, demand, and pricing creates a coupled decision-making problem, requiring models that reflect both the operator's strategic objectives and users' economic and behavioural responses.

Recent research has explored this integration through different methodological perspectives. Baffo et al. [7] examined EV charging planning from the viewpoint of local authorities, focusing on the optimal number, type, and spatial distribution of stations under a homogeneous pricing scheme. Their model implicitly assumes a monopolistic or centrally coordinated market structure, where a uniform price applies across the network. Although pricing is not explicitly optimized, their work highlights the importance of consistent price signals for ensuring equitable service provision and demand coverage.

A shift toward explicitly modelling pricing decisions can be seen in more recent studies. Adil et al. [1] adopted a bilevel Stackelberg framework to jointly optimize prices and station locations. In their model, the central planner acts as the leader, selecting infrastructure configurations and energy prices with the objective of maximizing social welfare or operator profit. Users, modeled as followers, respond to these decisions by choosing charging stations that minimize their generalized travel and charging cost. This hierarchical formulation effectively captures the feedback loop between pricing strategies and user behaviour, and highlights the potential of coordinated planning in enhancing both network performance and economic viability.

Zavvos et al. [59] studied pricing in a decentralized competitive environment with multiple investors aiming to maximize their own profits. Their game-theoretic model reflects market fragmentation, where each investor independently selects locations and pricing strategies while anticipating competitors' actions. This approach reveals how competition can lead to differentiated pricing structures, local price wars, and strategic clustering or dispersion of charging stations. It also shows how investor behaviour depends on expected demand elasticity, operational costs, and spatial characteristics of the service area.

1.2 Research questions

The increasing penetration of electric vehicles (EVs) and the parallel need for a robust, accessible, and economically sustainable charging infrastructure raise a set of strategic questions for charge point operators (CPOs), urban planners, and policymakers. Despite the growing body of literature on flow-based location models, refueling feasibility, deviation behaviour, and capacity-constrained planning, several critical challenges remain unresolved when real-world behavioural, spatial, and economic factors are incorporated. Building on the methodological and empirical research developed throughout this thesis—including machine-learning analysis of charging demand, graph-based modelling of spatial interactions, and the formulation of advanced optimization frameworks—this work aims to address the following research questions.

(1) How can a data-driven methodology be developed to identify profitable locations for EV charging stations?

Profitability in EV charging infrastructure is not determined solely by traffic volumes or coverage potential. It depends on a complex interplay of variables, including mobility patterns, the spatial structure of cities, proximity to relevant points of interest (POIs), expected utilization rates, local competition, and pricing strategies. This thesis investigates how these factors jointly influence station performance by integrating flow-based optimization models with data-driven demand prediction approaches. The objective is to derive a robust methodology for identifying high-potential sites that balance accessibility, demand feasibility, installation cost, grid constraints, and long-term revenue prospects.

(2) How can we support companies and public administrators in deciding where to build new charging stations?

The second research question focuses on translating analytical models into actionable decision-support tools. The thesis develops a modelling framework that combines classical location models (flow-capturing, flow-refueling, deviation-flow) with real network data and machine-learning–based demand estimation. This framework is capable of capturing key practical complexities—including charging duration, connector types, heterogeneous power levels, station capacity, user deviation behaviour, and demand uncertainty. By integrating these components, the research aims to provide companies with a systematic methodology for site selection that is both theoretically grounded and operationally viable, facilitating informed infrastructure expansion strategies.

(3) Which points of interest influence charging-station usage the most, and how can this information be incorporated into optimization models?

Understanding the relationship between the urban environment and charging behaviour is crucial for accurate demand modelling. Using real charging-session data complemented by POI information extracted from OpenStreetMap (OSM), this thesis investigates the extent to which POI categories—such as retail areas, offices, leisure facilities, residential density, or transport hubs—affect station utilization. Advanced machine-learning models, including random forests, XGBoost, and graph neural networks with attention mechanisms, are employed to uncover the relative importance and spatial interactions of POIs.

The insights gained are subsequently incorporated into the optimization framework in two ways: (i) as demand predictors that feed into location models, improving demand estimation accuracy, and (ii) as structural features that inform the attractiveness term in bilevel and user-behavioural models. In this way, the thesis bridges the gap between data-driven demand characterization and mathematical optimization for EV infrastructure planning.

(4) At what price should electricity for EV charging be sold to the public?

Pricing is a fundamental dimension of charging-station profitability and user choice. While classical location models typically assume exogenous prices, this thesis expands the modelling approach by incorporating pricing into a bilevel optimization structure. In the upper level, the operator sets prices with the objective of maximizing revenue or utilization subject to technical and capacity constraints. In the lower level, users choose charging stations based on generalized cost, which includes price sensitivity, deviation distance, connector power, and the attractiveness of surrounding POIs.

This bilevel formulation captures key economic and behavioural mechanisms influencing users' charging decisions, enabling a more realistic evaluation of pricing strategies. The research also explores feasible bounds on prices emerging from market competition, willingness to pay, and spatial attractiveness, offering practical insights for operators seeking to set prices dynamically while ensuring service accessibility.

1.3 Thesis structure

This dissertation is structured as follows.

Chapter 2 introduces the theoretical foundations underlying this work. It reviews the relevant literature and methodological background related to electric vehicle charging infrastructure, demand modeling, machine learning techniques, and optimization frameworks, providing the conceptual tools required for the subsequent chapters.

Chapter 3 describes the data sources and data construction process. It presents the real-world dataset on electric vehicle charging stations in Italy provided by Atlante Srl, and details the procedures adopted to build a realistic and spatially consistent representation of the city of Milan. This includes data cleaning, validation, enrichment with external information, and the generation of additional features required for the analysis.

Chapter 4 is devoted to preliminary data analysis. It explores the main characteristics of the datasets, investigates spatial and temporal patterns of charging activity, and highlights key relationships between infrastructure, mobility, and contextual variables. These analyses provide the empirical motivation for the predictive and optimization models developed later in the thesis.

Chapter 5 focuses on the predictive modeling of charging station utilization using machine learning techniques. Several models are developed and compared to estimate charging demand and station performance, with particular attention paid to feature relevance, model interpretability, and predictive uncertainty.

Chapter 6 presents the optimization framework proposed in this thesis. It introduces a mathematical formulation that integrates demand-related insights from the machine learning models into a location and design problem for electric vehicle charging stations, accounting for capacity constraints, pricing decisions, and user behavior. Some of the concepts and modeling choices discussed in this chapter build upon results previously presented in [21].

Finally, Chapter 7 concludes the thesis by summarizing the main findings, discussing their implications for planning and operation of electric vehicle charging infrastructure, and outlining possible directions for future research.

Chapter 2

Theory foundations

2.1 Linear programming basics

The study of finding optimal solutions, i.e., solutions that are at least as good as any other potential solution, is known as **mathematical optimization**. The solutions considered are generally elements of an n-dimensional vector space $\mathbb{R}^n$, that is, vectors with finitely many real components. Each component is regarded as a *decision variable*, in the sense that it encodes a specific decision to be made. The quality of a solution is measured through a function $f(x_1, \ldots, x_n)$ of the n decision variables, called the *objective function*, which must be maximized or minimized in a domain called the *feasible region*. The feasible region is generally defined implicitly by a system of inequalities that must be fulfilled by the decision variables, and a vector $\mathbf{x} \in \mathbb{R}^n$ that fulfills all these inequalities is said to be a *feasible solution*.

More formally, given a vector of decision variables $\mathbf{x} \in \mathbb{R}^n$, and a real-valued vector function $h : \mathbb{R}^n \rightarrow \mathbb{R}^m$, we define the feasible region as the subset of $\mathbb{R}^n$ formed by all those n-vectors $\mathbf{x}$ for which $h(\mathbf{x}) \leq \mathbf{0}$ holds true. We can, in this way, introduce a mathematical optimization problem in the form

$$\begin{aligned} \max \ f(x_1, \ldots, x_n) & \\ h_1(x_1, \ldots, x_n) &\leq 0 \\ &\cdots \\ h_m(x_1, \ldots, x_n) &\leq 0 \end{aligned}$$

If a solution $\mathbf{x}^* \in P$ meets all the constraints and achieves the global maximum of the objective function, it is said to be *optimal* (in case optimality is achieved by a minimum of function f, one can turn the problem in the form above by simply changing f into $-f$). A noticeable case is when $f(\mathbf{x}) = c_1x_1 + \ldots + c_nx_n = \mathbf{cx}$ and h is a linear operator, for example

$$h_i(x_1, \ldots, x_n) = a_{i1}x_1 + \ldots + a_{in}x_n - b_i$$

with $a_{ij}, b_i \in \mathbb{R}$ for $i = 1, \dots, m$ and $j = 1, \dots, n$. In this case, we have a *linear optimization problem*, that can be put more concisely in the general form

$$\begin{aligned} \max \ & \mathbf{cx} \\ & \mathbf{Ax} \leq b \end{aligned}$$

where $\mathbf{A} = \{a_{ij}\}$ is a real $m \times n$ matrix and $\mathbf{b} = \{b_i\}$ is a real m-vector. The geometric counterpart of the set $P = \{x \in \mathbb{R}^n | \mathbf{Ax} \leq \mathbf{b}\}$ is a *convex polyhedron*. With simple algebraic manipulation, one can always put every linear optimization problem in an equivalent form called *standard*:

$$\begin{aligned} \max \mathbf{cx} & \\ \mathbf{Ax} &= \mathbf{b} \\ \mathbf{x} &\geq \mathbf{0} \end{aligned}$$

Also the feasible region of a standard form is a convex polyhedron, since one can replace $\mathbf{Ax} = \mathbf{b}, \mathbf{x} \geq \mathbf{0}$ with the system of inequalities $\mathbf{Ax} \leq \mathbf{b}, -\mathbf{Ax} \leq -\mathbf{b}, -\mathbf{x} \leq \mathbf{0}$.

Like any general optimization problem, a linear optimization problem $\mathcal{P}$ can be

- *infeasible*: $\mathcal{P}$ has no solution, that is, $P = \emptyset$;
- *unbounded*: $\mathcal{P}$ has a solution, but has no optimal solution because, for any $\mathbf{x} \in P$, one can find a $\mathbf{y} \in P$ such that $\mathbf{cy} > \mathbf{cx}$;
- *with a finite optimum*: there exists an $\mathbf{x} \in P$ such that $\mathbf{cx} \geq \mathbf{cy}$ for any $\mathbf{y} \in P$.

One can show that, in the last case, an optimal solution to a linear programming problem $\mathcal{P}$ can always be found in a "corner" of the associated polyhedron P; this point, called *vertex*, complies with the following definition:

Definition 2.1. *Let P be a convex polyhedron. A vector $\mathbf{x} \in P$ is a* vertex *of P if there exists a $\mathbf{c} \in \mathbb{R}^n$ such that $\mathbf{cx} < \mathbf{cy}$ for all $\mathbf{y} \in P$ such that $\mathbf{y} \neq \mathbf{x}$.*

The definition of vertex is interchangeable with that of *extreme point*:

Definition 2.2. *Let P be a convex polyhedron. A vector $\mathbf{x} \in P$ is an* extreme point *of P if one cannot find $\mathbf{y}, \mathbf{z} \in P$, both different from $\mathbf{x}$, and a scalar $\lambda \in [0, 1]$, such that $\mathbf{x} = \lambda\mathbf{y} + (1 - \lambda)\mathbf{z}$.*

Definition 2.3. *Let $S = \{\mathbf{x}^1, \dots, \mathbf{x}^k\} \subseteq \mathbb{R}^n$ and let $\lambda_1, \dots, \lambda_k$ be nonnegative scalars such that $\sum_{i=1}^k \lambda_i = 1$.*

1. *The vector $\mathbf{y} = \sum_{i=1}^k \lambda_i \mathbf{x}^i$ is said to be a* convex combination *of the vectors in S;*

2. *The* convex hull *of S, indicated as conv(S), is the set of all the vectors* $\mathbf{y}$ *obtained as convex combinations of the vectors in S; one can prove that conv(S) is the (unique) minimal convex set containing S.*

Consider a polyhedron $P \subseteq \mathbb{R}^n$ defined in terms of linear constraints (equations and/or inequalities)

$$\begin{aligned} \mathbf{a}_i'\mathbf{x} &\geq b_i, & i &\in I_1 \\ \mathbf{a}_i'\mathbf{x} &\leq b_i, & i &\in I_2 \\ \mathbf{a}_i'\mathbf{x} &= b_i, & i &\in I_3 \end{aligned}$$

where $I_1,\ I_2,\ I_3$ are finite index sets, $\mathbf{a}_i'$ is a vector of $\mathbb{R}^n$ and b_i is a scalar,

Definition 2.4. *If $\mathbf{a}_i'\mathbf{x}^* = b_i$ for some $\mathbf{x}^* \in \mathbb{R}^n$ and some $i \in I_1 \cup I_2 \cup I_3$, we say that the i-th constraint is* active *at $\mathbf{x}^*$.*

If n constraints are active at a vector $\mathbf{x}^*$, then x^* is a solution of a system of n linear equations in n variables. This system has a unique solution if ad only if its equations are *linearly independent.* This leads to define a vertex as a feasible solution at which there are n linearly independent active constraints of P.

Definition 2.5. *Let $P \subseteq \mathbb{R}^n$ be a polyhedron defined in terms of linear equality and inequality constraints, and let $\mathbf{x}^* \in \mathbb{R}^n$*

1. $\mathbf{x}^*$ *is said to be a* basic solution *if*

   (a) all the equality constraints of P are active at $\mathbf{x}^*$;

   (b) n out of the active constraints at $\mathbf{x}^*$ are linearly independent.

2. *A basic solution that fulfills all the constraints that define P is called a* basic feasible solution.

Definitions 2.1, 2.2 and 2.5 are all equivalent: the first two (vertex and extreme point) express a geometric property, while the last one (basic feasible solution) an algebraic property. This is resumed by the following theorem:

Theorem 2.1. *Let P be a convex polyhedron, and let $\mathbf{x}^* \in P$. Then the following are equivalent:*

1. $\mathbf{x}^*$ *is a vertex;*

2. $\mathbf{x}^*$ *is an extreme point;*

3. $\mathbf{x}^*$ *is a basic feasible solution.*

Let us now fix a vector $\mathbf{y}$ in a polyhedron $P = \{\mathbf{x} : \mathbf{A}\mathbf{x} \geq \mathbf{b}\}$. We define the *recession cone* $\text{rec}(P)$ at $\mathbf{y}$ as:

$$\{\mathbf{d} \in \mathbb{R}^n \mid \mathbf{A}(\mathbf{y} + \lambda\mathbf{d}) \geq \mathbf{b}, \quad \forall\lambda \geq 0\}$$

or, more informally, as the set of all directions along which we can move indefinitely away from $\mathbf{y}$ without leaving the set P. The recession cone is always non-empty as it contains the origin $\mathbf{0}$: we have $\text{rec}(P) = \{0\}$ if and only if P is bounded (that is, is contained in a ball with finite radius). The nonzero elements of the recession cone, if any, are called the *rays* of P; if there are $n-1$ linearly independent constraints active at ray $\mathbf{u}$, then we call $\mathbf{u}$ an *extreme ray*. Every convex polyhedron has a finite number of extreme rays (possibly none).

One of the fundamental results of the Theory of Linear Programming states that any element of a convex polyhedron can be expressed as a convex combination of its extreme points plus a nonnegative linear combination of its extreme rays; formally:

Theorem 2.2. (Resolution theorem) *Let*

$$P = \{\mathbf{x} \in \mathbb{R}^n \mid \mathbf{A}\mathbf{x} \geq \mathbf{b}\}$$

be a non-empty polyhedron with at least one extreme point. Let $\mathbf{x}^1, \ldots, \mathbf{x}^{|Q|}$ be its extreme points and $\mathbf{u}^1, \ldots, \mathbf{u}^{|R|}$ be its extreme rays. Then, each $\mathbf{z} \in P$ can be expressed as

$$\mathbf{z} \;=\; \sum_{q \in Q} \alpha_q \mathbf{x}^q \;+\; \sum_{r \in R} \beta_r \mathbf{u}^r$$

with $\alpha_q, \beta_r \geq 0$, $\sum_{q \in Q} \alpha_q = 1$.

The *dual* $\mathcal{D}$ of a linear minimization problem $\mathcal{P}$ in standard form

$$\begin{aligned} \min\; & \mathbf{c}\mathbf{x} \\ & \mathbf{A}\mathbf{x} = \mathbf{b} \\ & \mathbf{x} \geq \mathbf{0} \end{aligned}$$

is a linear maximization problem defined as follows:

$$\begin{aligned} \max\; & \pi\mathbf{b} \\ & \pi\mathbf{A} \leq \mathbf{c} \end{aligned}$$

Reciprocally, $\mathcal{P}$ is called the *primal problem*. The dual bounds that can be achieved by linear programming are particularly valuable in practice and can provide a measure of the efficiency or effectiveness of the operational solution, regardless of the possibilities of finding an optimal solution.

This is because a solution of the dual problem provides a lower bound to the optimal solution of the primal.

Theorem 2.3. (Weak duality) *If* $\mathbf{x}$ *is a feasible solution of* $\mathcal{P}$ *and* $\mathbf{y}$ *is a feasible solution of* $\mathcal{D}$*, then*

$$\pi\mathbf{b} \leq \mathbf{cx}$$

Theorem 2.4. (Strong duality) *If* $\mathbf{x}^*$ *is an optimal solution of* $\mathcal{P}$*, then*

1. *there exists an optimal solution* π^* *of* $\mathcal{D}$ *and*
2. $\pi^*\mathbf{b} = \mathbf{cx}^*$

The *reduced cost* (or opportunity cost) of a nonbasic variable represents the marginal change in the objective function value associated with increasing that variable by one unit from its bound, while maintaining feasibility with respect to the current basis. Equivalently, it measures how much the objective coefficient of a nonbasic variable would need to improve —increase in a maximization problem or decrease in a minimization problem— for that variable to enter the optimal solution. According to the Theory of Linear Programming, a solution of a linear minimization problem in standard form is optimal if all the reduced costs are non-negative. Using the definition of dual problem, this amounts to say:

$$\mathbf{c} - \pi\mathbf{A} \geq \mathbf{0}$$

2.2 Dantzig-Wolfe decomposition

Consider a linear program (LP), usually called **compact formulation**

$$\min \quad \mathbf{cx} \tag{2.1}$$

$$\mathbf{Ax} \geq \mathbf{b} \tag{2.2}$$

$$\mathbf{Dx} \geq \mathbf{d} \tag{2.3}$$

$$\mathbf{x} \geq \mathbf{0} \tag{2.4}$$

Theorem 2.2 can be used to reformulate this linear program in an analogous manner. Notice first that the constraints are divided into two types: we will refer to (2.3) as the "easy constraints", which are the ones we can deal well with; while we will regard (2.2) as the "complicated constraints".

The idea for solving LP, or estimating its optimal value, is to exploit our ability to optimize over the polyhedron $P = \{\mathbf{x} \in \mathbb{R}^n \mid \mathbf{Dx} \geq \mathbf{d}\}$ defined by (2.3), the "easy constraints", which we assume to be non-empty.

Applying Theorem 2.2, we can substitute $\mathbf{x} \in P$ and rewrite formulation of LP in the equivalent extensive formulation (D_l):

$$\min \sum_{q \in Q} c_q \alpha_q + \sum_{r \in R} c_r \beta_r \tag{2.5}$$

$$\sum_{q \in Q} a_q \alpha_q + \sum_{r \in R} a_r \beta_r \geq b \tag{2.6}$$

$$\sum_{q \in Q} \alpha_q = 1 \tag{2.7}$$

$$\alpha_q \geq 0 \qquad q \in Q \tag{2.8}$$

$$\beta_r \geq 0 \qquad r \in R \tag{2.9}$$

where $c_q = \mathbf{c}\mathbf{x}^q$ and $c_r = \mathbf{c}\mathbf{u}^r$, and $\mathbf{x}^q, \mathbf{u}^r$ are defined as in Theorem 2.2. The equation $\sum_{q \in Q} \alpha_q = 1$ is referred to as the *convexity constraint*.

D_l has fewer rows than the compact formulation, but, due to the large number of extreme points and rays, it has a large number of variables, i.e., $(|Q| + |R|)$, that may grow exponentially making resolution impractical. As a result, most D_l problems are solved using a column generation technique that considers only a subset of columns, with missing columns generated as needed. Details on the column generation algorithm will be provided later.

2.2.1 Decomposition of Integer Programs

In most cases, we are interested in optimizing over a discrete set X, that is

$$z^* := \min \quad \mathbf{c}\mathbf{x} \tag{2.10}$$

$$\mathbf{A}\mathbf{x} \geq \mathbf{b} \tag{2.11}$$

$$\mathbf{x} \in X \tag{2.12}$$

where $X = P \cap \mathbb{Z}^+$ and $P \subseteq \mathbb{R}^n$ is a polyhedron.

Decomposition can be done using one of two methods when reformulating an *integer program* (IP) or a *mixed integer program* (MIP): convexification or discretization.

Convexification is achieved by substituting X with its *convex hull* $\text{conv}(X)$ and expressing each $\mathbf{x}$ in terms of $\text{conv}(X)$ extreme points and rays. However, we must place extra requirements on the x-variables in order to obtain integer solutions. These requirements will appear in the compact formulation at the master problem, subproblem, or both levels. In fact, since the optimum integer solution of (2.10)-(2.12) may be an interior point of $\text{conv}(X)$, requiring integrality of the variables of MPs does not result in an integer program equal to (2.10)-(2.12).

Alternatively, **discretization** can be seen as the integer counterpart of the decomposition principle. It is based on the following:

Theorem 2.5. *Let*

$$P = \{\mathbf{x} \in \mathbb{R}^n \mid \mathbf{D}\mathbf{x} \geq \mathbf{d},\ \mathbf{x} \geq \mathbf{0}\}$$

be a non-empty polyhedron with at least one extreme point and $X = P \cap \mathbb{Z}^+$. *Let* $\{\mathbf{x}^1,\ldots,\mathbf{x}^{|Q|}\} \subseteq X$ *be the integer points and* $\mathbf{u}^1,\ldots,\mathbf{u}^{|R|}$ *be the integer rays of* P*. Then, each* $\mathbf{x} \in X$ *can be expressed as*

$$\mathbf{x} \ = \ \sum_{q \in Q} \alpha_q \mathbf{x}^q \ + \ \sum_{r \in R} \beta_r \mathbf{u}^r$$

with $\alpha \in \mathbb{Z}^+_{|Q|},\ \beta \in \mathbb{Z}^+_{|R|},\ \sum_{q \in Q} \alpha_q = 1.$

So, Theorem 2.5 allows us to express each point of a polyhedron as a convex combination of integer points plus a non-negative combination of integer rays.

Substitution for $\mathbf{x}$ in (2.10)-(2.12) as given by Theorem 2.5 yields a MIP:

$$\min \sum_{q \in Q} c_q \alpha_q \ + \ \sum_{r \in R} c_r \beta_r \tag{2.13}$$

$$\sum_{q \in Q} a_q \alpha_q \ + \ \sum_{r \in R} a_r \beta_r \geq b \tag{2.14}$$

$$\sum_{q \in Q} \alpha_q = 1 \tag{2.15}$$

$$\alpha_q \in \mathbb{Z}^+_{|Q|} \tag{2.16}$$

$$\beta_r \in \mathbb{Z}^+_{|R|} \tag{2.17}$$

where $c_q = \mathbf{c}\mathbf{x}^q$ and $c_r = \mathbf{c}\mathbf{u}^r$.

Convexification and discretization coincide in the particular case of $X \subseteq [0,1]^n$, since all integral points in the bounded set X are already vertices of conv(X).

2.2.2 Block Diagonal Structure

When using decomposition techniques, it is common for problems to exhibit a matrix coefficient D with a modular structure, as shown in 2.2.2. Let K denote the set of independent blocks that can be used to partition D. We can identify $|K|$ disjoint polyhedra

$$P^k = \{\mathbf{x}^k \in \mathbb{R}^n \mid \mathbf{D}^k \mathbf{x}^k \geq \mathbf{d}^k\}$$

such that $X^k = P^k \cap \mathbb{Z}^+$. So, (2.10)-(2.12) can be written as

$$\min \sum_{k \in K} \mathbf{c}^k \mathbf{x}^k \tag{2.18}$$

$$\sum_{k \in K} \mathbf{A}^k \mathbf{x}^k \geq b \tag{2.19}$$

$$\mathbf{x}^k \in X^k \qquad \forall k \in K \tag{2.20}$$

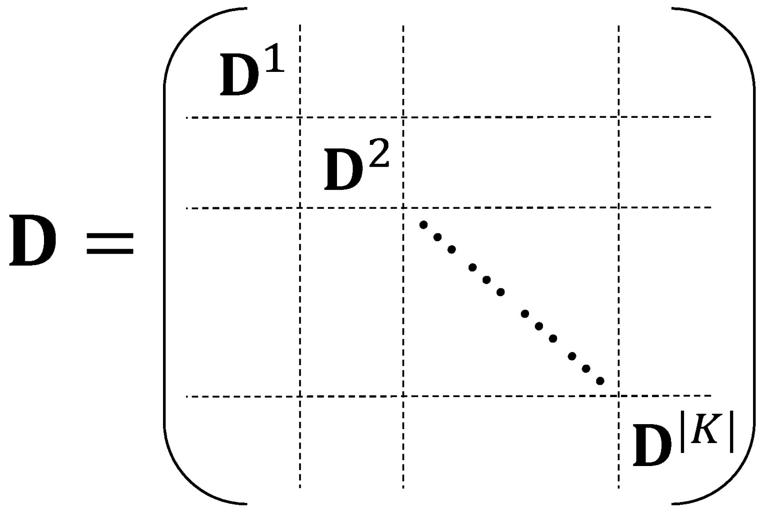

Figure 2.1: Block-diagonal structure

So, we can rewrite the MP as:

$$\min \sum_{k \in K} \sum_{q \in Q^k} c_q^k \alpha_q^k + \sum_{k \in K} \sum_{r \in R^k} c_r^k \beta_r^k \tag{2.21}$$

$$\sum_{k \in K} \sum_{q \in Q^k} a_q^k \alpha_q^k + \sum_{k \in K} \sum_{r \in R^k} a_r^k \beta_r^k \geq b \tag{2.22}$$

$$\sum_{q \in Q^k} \alpha_q^k = 1 \qquad \forall k \in K \tag{2.23}$$

$$\alpha^k \in \mathbb{Z}^+_{|Q|^k} \qquad \forall k \in K \tag{2.24}$$

$$\beta^k \in \mathbb{Z}^+_{|R|^k} \qquad \forall k \in K \tag{2.25}$$

A superscript k to c_q^k, c_r^k, α_q^k, β_r^k, indicates the respective subsystem $k \in K$. Applying the block structure translates into computational advantages, since this modular structure allows solving a set of disjoint pricing problems.

2.2.3 Lagrangian Relaxation

A popular approach to tackle (2.10)-(2.12) is *Lagrangian relaxation.* The core idea is to relax constraints $\mathbf{Ax} \geq \mathbf{b}$ and penalize their violation in the objective function via Lagrangian multipliers $\mathbf{u} \geq \mathbf{0}$. So, we obtain the following Lagrangian subproblem:

$$L(u) = \min_{\mathbf{x} \in X} \mathbf{cx} - \mathbf{u} \cdot (\mathbf{Ax} - \mathbf{b}). \tag{2.26}$$

$L(u)$ is a lower bound on z^*, because

$$L(u) \;\leq\; \min\{\mathbf{cx} - \mathbf{u} \cdot (\mathbf{Ax} - \mathbf{b}) \mid \mathbf{Ax} \geq \mathbf{b},\; \mathbf{x} \in X\} \;\leq\; z^*$$

The best of such bounds on z^* is computed in the *Lagrangian dual problem*:

$$\mathcal{L} := \max_{\mathbf{u} \geq \mathbf{0}} L(u). \tag{2.27}$$

Assume that we are given optimal multipliers $\mathbf{u}^*$ for (2.27). By solving (2.26), we ensure that $\mathbf{x} \in X$; the optimality of $\mathbf{u}^*$ implies that $\mathbf{u}^* \cdot (\mathbf{Ax} - \mathbf{b}) = 0$ (complementary slackness), but $\mathbf{Ax} \geq \mathbf{b}$ (feasibility) must be verified to prove optimality. If this condition is violated, the primal-dual pair $(\mathbf{x}, \mathbf{u}^*)$ is not optimal. Moreover, if X is replaced by $\operatorname{conv}(X)$ in (2.10)-(2.12), this does not change z^*.

2.3 Bilevel Optimization: Models, Reformulations, and Algorithms

Bilevel optimization models describe hierarchical decision-making processes in which a *leader* (upper level) anticipates the rational response of a *follower* (lower level). They provide a mathematical framework for modeling Stackelberg-type problems arising in economics, engineering, transportation, and energy systems [8, 16, 17, 54]. Formally, a standard bilevel program can be written as

$$\begin{aligned} \min_{\mathbf{x} \in X,\, \mathbf{y} \in Y} \quad & F(\mathbf{x}, \mathbf{y}) \\ \text{s.t.} \quad & G(\mathbf{x}, \mathbf{y}) \leq \mathbf{0}, \\ & \mathbf{y} \in \arg\min_{\mathbf{z} \in Y} \left\{ f(\mathbf{x}, \mathbf{z}) \;:\; g(\mathbf{x}, \mathbf{z}) \leq \mathbf{0} \right\}, \end{aligned} \tag{2.28}$$

where $\mathbf{x}$ denotes the leader's decision variables and $\mathbf{y}$ the follower's decision variables. The leader selects $\mathbf{x}$, anticipating that the follower will optimally react by solving the lower-level problem parameterized by $\mathbf{x}$.

2.3.1 Optimistic vs. Pessimistic Formulations and Solution Concepts

A central modeling choice in bilevel optimization concerns how to handle the presence of multiple optimal solutions at the follower level. In the *optimistic* formulation, the follower is assumed to resolve ties in a way that is most favorable to the leader, whereas in the *pessimistic* formulation, ties are resolved in the least favorable manner for the leader [16, 17].

The existence, uniqueness, and stability of bilevel solutions depend critically on regularity assumptions, such as compactness of the feasible sets, continuity of the objective and constraint functions, and the satisfaction of appropriate constraint qualifications, as well as on the adopted selection rule (optimistic or pessimistic). From the leader's perspective, bilevel problems typically induce nonconvex and possibly disconnected feasible regions, even when both the upper- and lower-level problems are linear. As a consequence, bilevel optimization problems are computationally challenging and are, in general, NP-hard; moreover, several important subclasses are known to be strongly NP-hard [11, 14].

2.3.2 Important Model Classes

This section briefly reviews the main classes of bilevel optimization models that are most commonly encountered in the literature, highlighting their key properties and computational implications.

Bilevel Linear Programs (BLP). In bilevel linear programs, both the upper- and lower-level problems are linear. Despite their apparent simplicity, BLPs exhibit intrinsically nonconvex feasible regions for the leader and are known to be NP-hard. These models arise frequently in applications such as pricing problems, network interdiction, and network design [14].

Convex Lower-Level Problems (BLO). In this class of models, the follower's problem is convex for each fixed value of the leader's decision variables, and is often assumed to be strictly or strongly convex. Under suitable constraint qualifications, such as the Linear Independence Constraint Qualification (LICQ), and assuming that strong duality holds, the optimality conditions of the lower-level problem can be characterized via the Karush–Kuhn–Tucker (KKT) conditions. This property enables reformulations of the bilevel problem as a single-level Mathematical Program with Equilibrium Constraints (MPEC), which can then be tackled using specialized nonlinear or mixed-integer optimization techniques [16, 36].

Mixed-Integer Bilevel Problems. Mixed-integer bilevel optimization models arise when integrality constraints are imposed at one or both levels, leading to Mixed-Integer

Bilevel Linear Programs (MIBLP) or Mixed-Integer Linear Bilevel problems (MILB). These formulations are particularly relevant in applications such as power market modeling, network design, and facility location. Due to the combination of hierarchical structure, nonconvexity, and discrete decisions, these problems are especially challenging. State-of-the-art solution approaches typically rely on decomposition schemes, branch-and-bound strategies, and cutting-plane methods [28, 60].

2.3.3 Single-Level Reformulations via KKT Conditions

Several approaches have been proposed in the literature to transform the bilevel problem in (2.28) into an equivalent (or approximate) single-level formulation, which is generally nonconvex. Two classical reformulation strategies are briefly reviewed below.

KKT-Based (MPEC) Reformulation. When the follower's problem is convex for a given leader decision $\mathbf{x}$ and satisfies strong duality, the set-valued optimality condition $\mathbf{y} \in \arg\min\{f(\mathbf{x},\mathbf{z}) : g(\mathbf{x},\mathbf{z}) \leq \mathbf{0}\}$ can be equivalently characterized by the Karush–Kuhn–Tucker (KKT) conditions. Embedding the follower's stationarity, primal feasibility, dual feasibility, and complementarity conditions into the upper-level problem leads to a *mathematical program with equilibrium constraints* (MPEC) [36, 40].

The resulting formulation is typically nonconvex due to the presence of complementarity constraints. In linear or quadratic settings, such complementarity relations can be reformulated using big-M techniques or special ordered sets of type 1 (SOS1), provided that appropriate bounds on the variables are available, yielding mixed-integer linear or nonlinear programs. In nonlinear contexts, alternative approaches include exact penalty methods, regularization techniques, and smoothing schemes, which aim to approximate or relax the complementarity structure while preserving tractability [17, 19, 40].

Value-Function Reformulation. An alternative reformulation strategy relies on the explicit use of the follower's value function. In this case, the lower-level optimality condition is enforced by imposing $f(\mathbf{x},\mathbf{y}) \leq \varphi(\mathbf{x})$, where the value function $\varphi(\mathbf{x})$ is defined as $\varphi(\mathbf{x}) = \min\{f(\mathbf{x},\mathbf{z}) : g(\mathbf{x},\mathbf{z}) \leq \mathbf{0}\}$.

This approach avoids the explicit introduction of complementarity constraints, but replaces them with the value function, which is generally nonsmooth and difficult to evaluate. When the follower's problem is convex, tools from convex analysis and subdifferential calculus can be employed to characterize and approximate $\varphi(\mathbf{x})$, enabling theoretical analysis and the development of solution algorithms [16, 17].

2.3.4 Primal–Dual Reformulation via Strong Duality

When the follower problem is convex and satisfies strong duality, bilevel programs can be reformulated by explicitly introducing both the *primal* and the *dual* formulations of the lower-level problem. This approach, originally formalized in the bilevel optimization literature (see [16, 17]) and rooted in classical convex duality theory (e.g., [9]), provides an alternative to KKT-based MPEC reformulations, which typically involve nonconvex complementarity constraints.

The key idea is to exploit the fact that, under suitable regularity assumptions such as Slater's condition, the optimal value of the follower's primal problem coincides with that of its dual. This property allows the lower-level optimality condition

$$\mathbf{y} \in \arg\min_{\mathbf{z}} \{f(\mathbf{x}, \mathbf{z}) : g(\mathbf{x}, \mathbf{z}) \leq \mathbf{0}\}$$

to be replaced by primal feasibility, dual feasibility, and a *strong duality equation*, following the classical framework discussed in [36, 40].

Setting

Consider a bilevel model in which the follower solves the convex optimization problem

$$\begin{aligned} \min_{\mathbf{z} \in \mathbb{R}^m} \quad & f(\mathbf{x}, \mathbf{z}) \\ \text{s.t.} \quad & g_i(\mathbf{x}, \mathbf{z}) \leq 0, \quad i = 1, \ldots, p \end{aligned} \tag{2.29}$$

where $f(\cdot)$ and $g_i(\cdot)$ are convex functions of $\mathbf{z}$ for any fixed $\mathbf{x}$. Assuming that Slater's condition holds, the corresponding dual problem is well defined and attains its optimum [9]:

$$\max_{\lambda \geq 0} \quad d(\mathbf{x}, \lambda) := \inf_{\mathbf{z}} \Big\{ f(\mathbf{x}, \mathbf{z}) + \sum_{i=1}^{p} \lambda_i g_i(\mathbf{x}, \mathbf{z}) \Big\}. \tag{2.30}$$

Strong Duality

For convex problems satisfying Slater's condition, the strong duality theorem [9] ensures that

$$f(\mathbf{x}, \mathbf{y}) = d(\mathbf{x}, \lambda^\star), \tag{2.31}$$

where $\lambda^\star$ denotes an optimal dual solution. As emphasized in [16, 17], this equality replaces the complementarity and stationarity conditions required in KKT-based MPEC reformulations [36].

Explicit Primal–Dual System

Under the above assumptions, follower optimality can be fully characterized by the following system:

$$\begin{aligned} & g_i(\mathbf{x}, \mathbf{y}) \leq 0, && i = 1, \ldots, p && \text{(primal feasibility)}, \\ & \lambda_i \geq 0, && i = 1, \ldots, p && \text{(dual feasibility)}, \\ & f(\mathbf{x}, \mathbf{y}) = d(\mathbf{x}, \lambda), && \text{(strong duality)}. \end{aligned} \tag{2.32}$$

This formulation avoids the explicit complementarity constraints inherent to KKT-based approaches and often leads to more tractable single-level models, particularly in linear and convex quadratic settings [14, 28].

Linear Case

When the follower problem is linear, i.e.,

$$f(\mathbf{x}, \mathbf{z}) = \mathbf{c}(\mathbf{x}) \cdot \mathbf{z}, \qquad g(\mathbf{x}, \mathbf{z}) = \mathbf{A}(\mathbf{x}) \cdot \mathbf{z} - \mathbf{b}(\mathbf{x}),$$

the dual problem is also linear, and the primal–dual reformulation yields the system

$$\begin{aligned} \mathbf{A}(\mathbf{x}) \cdot \mathbf{y} - \mathbf{b}(\mathbf{x}) &\leq \mathbf{0} \\ \lambda &\geq \mathbf{0} \\ \mathbf{A}(\mathbf{x}) \cdot \lambda &= \mathbf{c}(\mathbf{x}) \\ \mathbf{c}(\mathbf{x}) \cdot \mathbf{y} &= \mathbf{b}(\mathbf{x}) \cdot \lambda \end{aligned} \tag{2.33}$$

which is a classical result in the bilevel optimization literature [14, 60]. This reformulation leads to a mixed-integer linear or quadratic program without complementarity constraints, often providing a significant computational advantage over MPEC formulations [40].

Advantages

The primal–dual approach:

- avoids complementarity constraints (a major source of nonconvexity in MPECs [36]),
- produces tighter and more stable single-level reformulations (as observed in [28]),
- exploits explicit dual structure, often yielding more compact models,
- is particularly efficient for linear and convex quadratic lower-level problems.

Limitations

The approach requires:

- convexity of the follower problem,
- validity of strong duality (Slater's condition),
- availability of a closed-form dual function.

These limitations are well documented in foundational monographs such as [16, 17].

2.3.5 Algorithms

Exact and Global Approaches. For BLP and certain BLO settings, exact global algorithms include specialized branch-and-bound/branch-and-cut on the leader variables with lower-level optimality cuts; cutting-plane and column-and-constraint generation (C&CG) schemes that iteratively add follower-optimality constraints or no-good cuts; and decomposition exploiting strong duality of the follower [17, 28, 60]. In linear bilevel interdiction, C&CG is particularly effective [60].

MPEC Solvers and Penalization. KKT-based single-level problems can be attacked by nonlinear programming with constraint relaxations (e.g., Scholtes or Kanzow-type penalizations), or by mixed-integer formulations using big-M/SOS1 to model complementarity [17, 40].

Gradient-Based Methods for Smooth BLO. When the follower is strongly convex and smooth, the follower solution mapping $y^\star(x)$ is single-valued and (under regularity) differentiable via the implicit function theorem. This enables *implicit differentiation* and bilevel chain rules to compute $\nabla_x F(x, y^\star(x))$ without solving KKT systems explicitly at each step [20, 42]. Such methods underpin modern applications in hyperparameter optimization, meta-learning, and machine learning [20].

Heuristics and Evolutionary Methods. Evolutionary/metaheuristic approaches (GA, DE, PSO) have been widely explored for general, nonconvex/non-smooth bilevel problems where exact guarantees are difficult [49]. They trade optimality certificates for robustness and simplicity.

2.3.6 Modeling Choices and Practical Issues

Regularity and Constraint Qualifications. KKT-based reformulations require conditions such as LICQ and strong duality at the follower level; otherwise, MPEC stationarity may be insufficient and spurious stationary points may arise [17, 36].

Big-M Calibration. MILP/MINLP encodings of complementarity need bounding constants. Tighter problem-specific bounds improve numerics and correctness (avoid weak relaxations and false feasibility) [28].

Optimistic vs. Pessimistic. The pessimistic case is often harder; additional disjunctions or adversarial selection variables are required to represent worst-case follower choices [17].

Stochastic and Robust Bilevel. Uncertainty in data (demands, prices) motivates stochastic/robust variants, where either level hedges against scenarios or uncertainty sets. Solution techniques extend decomposition with scenario-wise cuts or robust counterparts (see surveys [14, 28]).

2.3.7 Applications (Selected)

Bilevel models naturally capture *leader–follower* interactions: (i) network pricing and toll-setting; (ii) facility location with user equilibrium routing; (iii) interdiction and fortification; (iv) electricity markets and transmission expansion with market clearing; (v) learning and hyperparameter tuning in machine learning [8, 14, 17, 20]. In charging-infrastructure planning, bilevel models appear when an operator selects locations, capacities, or prices anticipating user route/charge choices at the lower level (often a convex or equilibrium model), enabling a principled treatment of induced demand and competition.

2.4 Machine Learning Models for Charging Demand Prediction

The study of electric mobility, and in particular the prediction of demand for charging infrastructure, requires models that can capture complex, nonlinear relationships between a wide range of explanatory factors. These include geographic, social, and traffic-related attributes, as well as the spatial distribution of points of interest. Classical statistical approaches are often too restrictive to represent such high-dimensional and heterogeneous interactions. For this reason, machine learning methods have become essential tools in the analysis and forecasting of charging demand.

This section introduces two supervised learning approaches that are particularly relevant in this context: *Random Forests* and *Extreme Gradient Boosting (XGBoost)*. Random Forests are ensemble models based on bagging, well-suited for handling heterogeneous features and providing robust predictions with limited parameter tuning. XGBoost represents a more advanced boosting method that constructs trees sequentially, correcting residual errors while incorporating regularization and second-order optimization; it is widely adopted for tabular data due to its high predictive accuracy.

To evaluate the predictive quality of these models in the context of charging station demand estimation, two metrics are employed: the coefficient of determination (R^2), which quantifies the proportion of variability in charging demand explained by the model, and the Root Mean Squared Error (RMSE), which expresses the average prediction error in the same units as the observed charging sessions. These complementary measures allow us to assess both explanatory power and practical predictive accuracy.

The following subsections provide a detailed overview of these machine learning models and performance metrics, establishing the methodological foundations for their application in the analysis of charging infrastructure demand.

2.4.1 Correlation Coefficients

Correlation coefficients are statistical measures that quantify the strength and direction of the relationship between two variables. Here, we describe three commonly used correlation coefficients: Pearson's r [41], Spearman's ρ [50], and Kendall's τ [25].

Pearson's r (Pearson Product-Moment Correlation Coefficient)

Pearson's r measures the linear relationship between two variables. It is defined as the covariance of the variables divided by the product of their standard deviations. The formula for Pearson's r is given by:

$$r = \frac{\sum_{i=1}^{n}(X_i - \overline{X})(Y_i - \overline{Y})}{\sqrt{\sum_{i=1}^{n}(X_i - \overline{X})^2}\sqrt{\sum_{i=1}^{n}(Y_i - \overline{Y})^2}} \tag{2.34}$$

where

- r is Pearson's correlation coefficient,
- X_i and Y_i are the individual sample points indexed with i,
- $\overline{X}$ and $\overline{Y}$ are the means of the X and Y variables, respectively,
- n is the number of sample points.

Spearman's ρ (Spearman's Rank Correlation Coefficient)

Spearman's ρ measures the strength and direction of the association between two ranked variables. It is often used when the data do not meet the normality assumption required for Pearson's r. The formula for Spearman's ρ is:

$$\rho = 1 - \frac{6\sum d_i^2}{n(n^2 - 1)} \tag{2.35}$$

where

- ρ is Spearman's rank correlation coefficient,
- d_i is the difference between the ranks of corresponding values X_i and Y_i,
- n is the number of observations.

Kendall's τ (Kendall Rank Correlation Coefficient)

Kendall's τ is a measure of the correspondence between two rankings and is used to assess the ordinal association between two measured quantities. The formula for Kendall's τ is:

$$\tau = \frac{2}{n(n-1)} \sum_{i<j} \text{sgn}(x_i - x_j)\text{sgn}(y_i - y_j) \tag{2.36}$$

where

- τ is Kendall's tau coefficient,
- n is the number of observations,
- $\text{sgn}(x_i - x_j)$ and $\text{sgn}(y_i - y_j)$ are the sign functions of the differences between the pairs of observations.

2.4.2 Random Forest Regressor

Random Forest (RF) is an ensemble learning method introduced by Breiman [10], which combines the predictions of multiple decision trees in order to improve predictive performance and reduce overfitting. While originally proposed for classification tasks, the method has also been widely applied to regression problems, where it is commonly referred to as the *Random Forest Regressor*.

Decision Trees as Base Learners

The fundamental building block of a Random Forest is the decision tree. A decision tree recursively partitions the feature space into disjoint regions by applying a sequence of splitting rules based on input variables. At each internal node, the feature and the split threshold are chosen to maximize a purity criterion. For regression tasks, the most common criterion is the reduction of variance, where the impurity of a node t is defined as

$$I(t) = \frac{1}{N_t} \sum_{i \in t} (y_i - \bar{y}_t)^2 , \tag{2.37}$$

with N_t denoting the number of samples in node t, y_i the target value of sample i, and $\bar{y}_t$ the mean target value of node t.

A regression tree produces a piecewise-constant approximation of the target function, where each terminal node (leaf) predicts the average of the target values of the training samples contained in that leaf. However, individual trees are highly prone to overfitting, especially when grown to large depths, and are sensitive to small perturbations in the data.

Ensemble Construction via Bagging

Random Forests address the instability of individual trees by constructing an ensemble through *bootstrap aggregating* (bagging). Given a training set $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^N$, the algorithm generates B bootstrap samples $\mathcal{D}^{(1)}, \ldots, \mathcal{D}^{(B)}$ by sampling N instances with replacement from $\mathcal{D}$. For each bootstrap sample $\mathcal{D}^{(b)}$, a decision tree $h^{(b)}(\mathbf{x})$ is trained.

The prediction of the Random Forest Regressor for a new instance $\mathbf{x}$ is the average of the predictions from the individual trees:

$$\hat{f}_{\mathrm{RF}}(\mathbf{x}) = \frac{1}{B} \sum_{b=1}^{B} h^{(b)}(\mathbf{x}). \tag{2.38}$$

By averaging, the variance of the prediction is reduced compared to a single tree, while the bias remains approximately the same.

Random Feature Selection

To further decorrelate the base learners and reduce variance, Random Forests introduce an additional layer of randomness by selecting, at each split of a tree, a random subset of m features from the full set of p available features ($m < p$). The best split is then chosen only among these m features. This procedure ensures that not all trees rely on the same dominant predictors, thus enhancing diversity within the ensemble.

Bias-Variance Trade-off

The predictive performance of the Random Forest can be understood through the bias-variance decomposition. A single deep tree has low bias but high variance. By aggregating multiple trees trained on bootstrapped samples and random subsets of features, the Random Forest substantially reduces variance while preserving low bias. Consequently, Random Forests typically achieve high predictive accuracy with limited overfitting.

Feature Importance

An additional advantage of Random Forests is their ability to quantify the relative importance of features. Feature importance can be computed by measuring the total decrease in impurity attributable to each feature across all trees, normalized over the forest. More recent approaches also consider permutation importance, where the predictive accuracy is measured before and after randomly permuting the values of a feature, thus evaluating the extent to which the feature contributes to predictive performance.

Advantages and Limitations

Random Forest Regressors present several advantages:

- They are robust to overfitting, thanks to ensemble averaging.

- They can handle high-dimensional datasets and complex, non-linear relationships.
- They provide feature importance measures, useful for model interpretation.
- They require limited parameter tuning, with the number of trees B and the number of features m being the main hyperparameters.

Nonetheless, Random Forests also exhibit some limitations:

- Predictions are less interpretable compared to simple decision trees.
- Training can be computationally expensive for very large datasets and a high number of trees.
- Since predictions are obtained by averaging, they may be less accurate for extrapolation outside the range of the training data.

2.4.3 Extreme Gradient Boosting (XGBoost)

Extreme Gradient Boosting (XGBoost) is a scalable and efficient implementation of the gradient boosting framework, introduced by Chen and Guestrin [12]. It has become one of the most widely used machine learning methods due to its predictive performance, computational efficiency, and flexibility in handling a variety of tasks, including regression, classification, and ranking.

Gradient Boosting Framework

XGBoost is based on the principle of gradient boosting, where an ensemble of weak learners (typically decision trees) is constructed in a stage-wise manner. Unlike bagging-based methods such as Random Forests, boosting focuses on sequentially improving the model by fitting new learners to the residual errors of the current ensemble.

Formally, given a training dataset $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^{N}$ with input vectors $\mathbf{x}_i \in \mathbb{R}^p$ and target values $y_i \in \mathbb{R}$, the objective function at iteration t is defined as

$$\mathcal{L}^{(t)} = \sum_{i=1}^{N} l\big(y_i, \hat{y}_i^{(t-1)} + f_t(\mathbf{x}_i)\big) + \Omega(f_t), \tag{2.39}$$

where $l(\cdot)$ is a differentiable loss function (e.g., squared error for regression), $\hat{y}_i^{(t-1)}$ is the current prediction of the ensemble, f_t is the newly added function (tree), and $\Omega(f_t)$ is a regularization term penalizing the complexity of f_t.

The additive model is updated iteratively:

$$\hat{y}_i^{(t)} = \hat{y}_i^{(t-1)} + f_t(\mathbf{x}_i). \tag{2.40}$$

Second-Order Approximation

To efficiently optimize the objective function, XGBoost employs a second-order Taylor expansion of the loss around the current predictions:

$$\mathcal{L}^{(t)} \approx \sum_{i=1}^{N} \left[g_i f_t(\mathbf{x}_i) + \frac{1}{2} h_i f_t(\mathbf{x}_i)^2 \right] + \Omega(f_t), \tag{2.41}$$

where $g_i = \partial_{\hat{y}^{(t-1)}} l(y_i, \hat{y}_i^{(t-1)})$ and $h_i = \partial^2_{\hat{y}^{(t-1)}} l(y_i, \hat{y}_i^{(t-1)})$ are the first- and second-order gradients, respectively. This approximation allows XGBoost to exploit both first- and second-order information, improving convergence speed and stability compared to traditional gradient boosting methods that only use first-order gradients.

Regularized Objective

A key contribution of XGBoost is its explicit regularization of tree complexity, which helps to prevent overfitting. The regularization term is defined as

$$\Omega(f) = \gamma T + \frac{1}{2} \lambda \sum_{j=1}^{T} w_j^2, \tag{2.42}$$

where T is the number of leaves in the tree, w_j is the weight of leaf j, and $\gamma, \lambda \geq 0$ are regularization parameters controlling the complexity of the tree. This penalization encourages simpler trees and thus improves the generalization ability of the model.

Tree Structure Optimization

Each regression tree f_t is defined as a mapping from input space to leaf weights:

$$f_t(\mathbf{x}) = w_{q(\mathbf{x})}, \quad q : \mathbb{R}^p \to \{1, \ldots, T\}, \tag{2.43}$$

where $q(\cdot)$ assigns each instance to a leaf index, and $w_{q(\mathbf{x})}$ is the corresponding leaf weight. The optimal leaf weights can be obtained in closed form:

$$w_j^* = -\frac{\sum_{i \in I_j} g_i}{\sum_{i \in I_j} h_i + \lambda}, \tag{2.44}$$

where I_j is the set of training instances in leaf j. The gain of a split is then computed as the reduction in the regularized objective, enabling efficient greedy construction of trees.

Shrinkage and Subsampling

XGBoost incorporates additional mechanisms to enhance performance:

- **Shrinkage (learning rate):** A scaling factor $\eta \in (0, 1]$ is applied to each newly added tree, reducing its contribution to the ensemble:

$$\hat{y}_i^{(t)} = \hat{y}_i^{(t-1)} + \eta f_t(\mathbf{x}_i). \tag{2.45}$$

  This slows down learning and allows more trees to be added, improving generalization.

- **Subsampling:** Both instance subsampling (randomly sampling training data) and feature subsampling (randomly selecting a subset of features for each tree) are supported, further reducing variance and overfitting.

Computational Efficiency

XGBoost is designed with computational efficiency in mind. It leverages:

- A sparsity-aware algorithm capable of handling missing values and sparse feature matrices.
- Parallelization across cores and distributed computing for scalability.
- A novel weighted quantile sketch algorithm for efficient approximate tree learning.

These features allow XGBoost to scale to large datasets while maintaining high performance.

Advantages and Limitations

XGBoost offers several advantages:

- High predictive accuracy due to its use of second-order optimization and regularization.
- Robustness to overfitting, enabled by shrinkage, subsampling, and explicit regularization.
- Flexibility in handling diverse data types and loss functions.
- Computational efficiency and scalability to very large datasets.

Nevertheless, XGBoost also has some limitations:

- The model is less interpretable compared to linear models or single decision trees.
- Training can still be computationally demanding for extremely large parameter spaces or ensembles.
- Hyperparameter tuning is often required to achieve optimal performance, which can be costly.

2.4.4 Comparison between Random Forest and XGBoost

Random Forest and XGBoost are among the most popular ensemble methods for regression and classification, yet they are based on different ensemble learning paradigms. Their comparison can be understood along several dimensions.

Ensemble Strategy. Random Forest relies on *bagging* (bootstrap aggregating), where multiple decision trees are trained independently on different bootstrap samples and their predictions are averaged. This procedure reduces variance and improves robustness. In contrast, XGBoost is based on *boosting*, where trees are added sequentially and each new tree is trained to correct the residual errors of the current ensemble. This approach reduces bias while carefully controlling variance through regularization and shrinkage.

Bias-Variance Trade-off. Random Forest achieves variance reduction by averaging decorrelated trees, but it does not systematically reduce bias, since all trees are grown in parallel. Conversely, XGBoost builds trees sequentially, reducing both bias and variance, although at the cost of increased sensitivity to parameter tuning and the risk of overfitting if regularization is insufficient.

Regularization and Control of Complexity. XGBoost explicitly regularizes model complexity by penalizing the number of leaves and the magnitude of leaf weights. Additionally, shrinkage (learning rate) and subsampling strategies offer further control over model capacity. Random Forest, on the other hand, relies mainly on averaging, feature subsampling, and limiting tree depth as implicit regularization mechanisms, without an explicit regularization term in the objective function.

Interpretability. Both methods sacrifice interpretability compared to single decision trees. However, Random Forest is often considered more interpretable, as feature importance measures derived from impurity reduction or permutation tests are easier to compute and interpret. While XGBoost also provides feature importance measures, the sequential boosting process complicates interpretation, especially when many trees are involved.

Computational Efficiency. XGBoost is specifically engineered for computational efficiency, with support for parallelization, distributed training, and optimized handling of sparse data. While Random Forest can also be parallelized, its scalability and optimization techniques are generally less advanced compared to XGBoost. For very large-scale problems, XGBoost often outperforms Random Forest in both training time and predictive accuracy.

Use Cases. Random Forest is typically preferred when interpretability, robustness, and reduced hyperparameter tuning are important considerations. It performs well out-of-the-box with limited parameter adjustment. XGBoost is usually chosen for tasks requiring state-of-the-art accuracy and efficiency, particularly in competitive machine learning settings, although it demands more careful hyperparameter tuning.

2.4.5 Performance Metrics: R^2 and RMSE

The evaluation of regression models typically relies on statistical metrics that quantify the agreement between predicted and observed values. Among the most widely used measures are the coefficient of determination (R^2) and the Root Mean Squared Error (RMSE). These metrics provide complementary perspectives: while R^2 quantifies the proportion of explained variance, RMSE expresses the average magnitude of prediction errors in the same units as the target variable.

Coefficient of Determination (R^2)

The coefficient of determination, denoted by R^2, assesses the proportion of variance in the dependent variable that is explained by the model. Given true values $\{y_i\}_{i=1}^N$ and predictions $\{\hat{y}_i\}_{i=1}^N$, it is defined as

$$R^2 = 1 - \frac{\sum_{i=1}^N (y_i - \hat{y}_i)^2}{\sum_{i=1}^N (y_i - \bar{y})^2}, \tag{2.46}$$

where $\bar{y} = \frac{1}{N}\sum_{i=1}^N y_i$ is the sample mean of the observed values.

Interpretation.

- $R^2 = 1$ indicates perfect prediction, where the model captures all variability in the data.
- $R^2 = 0$ means the model performs no better than predicting the mean of the data.
- $R^2 < 0$ occurs when the model performs worse than the naive mean predictor.

Although widely adopted, R^2 does not measure the accuracy of predictions in absolute terms, nor does it indicate whether the model is unbiased.

Root Mean Squared Error (RMSE)

The Root Mean Squared Error (RMSE) is defined as the square root of the average squared prediction error:

$$\text{RMSE} = \sqrt{\frac{1}{N}\sum_{i=1}^N (y_i - \hat{y}_i)^2}. \tag{2.47}$$

Unlike R^2, the RMSE is expressed in the same units as the target variable, providing an interpretable measure of the average deviation between predictions and observations.

Interpretation.

- A lower RMSE indicates better predictive accuracy.
- RMSE is sensitive to large errors due to the squaring operation, making it particularly suitable in contexts where large deviations are undesirable.
- Unlike R^2, RMSE does not provide a relative measure against a baseline (such as the sample mean), but rather an absolute measure of error magnitude.

Complementarity of R^2 and RMSE

While both metrics derive from squared errors, they capture different aspects of model performance:

- R^2 is a relative measure of fit, indicating how well the model explains variability compared to a baseline.
- RMSE is an absolute measure, indicating the typical prediction error magnitude in the original data scale.

For this reason, reporting both R^2 and RMSE provides a more complete evaluation: R^2 assesses the explanatory power of the model, while RMSE quantifies its practical predictive accuracy.

Chapter 3

Data Description

Reliable and well-structured data are the cornerstone of any data-driven modeling effort. Whether the objective is to build accurate machine learning models or to develop robust optimisation frameworks, the quality, consistency, and interpretability of input data fundamentally influence the validity and applicability of the results. Inadequate preprocessing or overlooked inconsistencies can lead to biased conclusions, misleading insights, or poor model performance.

To ground the reader in the physical and operational complexity of the EV charging ecosystem, the chapter begins by outlining the structural components of a charging station. This includes a discussion on the types of connectors (such as Type 1, Type 2, CCS, and CHAdeMO), the classification of charging modes according to international standards (IEC 61851 and IEC 62196), and the roles of key stakeholders involved in the charging infrastructure—namely, the Charging Point Operators (CPOs), eMobility Service Providers (eMSPs), Roaming Platforms, and EV drivers. Understanding the layered architecture of this system is essential for interpreting how data is generated, aggregated, and managed across different platforms and actors.

Building on this foundational understanding, a significant portion of the chapter is then devoted to the meticulous construction of a reliable and enriched dataset. The primary source of data is a leading eRoaming platform that collects and distributes real-time data from CPOs across Europe. These data are complemented by a range of open datasets detailing Italian infrastructure, administrative boundaries, population statistics, and mobility patterns—allowing for a rich, multi-dimensional representation of the charging landscape. The analysis is geographically limited to Italy and focuses on data from the year 2023 to ensure temporal consistency.

The methodological pipeline involves multiple stages of cleaning, validation, and contextual enrichment. These include the detection and correction of geographic inaccuracies, the standardization of station identifiers, the verification of connector-level details, and the temporal tracking of commissioning and decommissioning events. Stations are aggregated from connector-level data, and infrastructural dynamics such as rebranding, relocation, or upgrades are systematically identified. Each observation is also enriched with relevant external attributes—including regional traffic estimates, Points of Interest (POIs) within walking and driving distances, and local population densities.

The rest of the chapter is organized as follows: the second section of the chapter describes the multi-step data validation and cleaning procedures designed to ensure the accuracy and coherence of the dataset. These procedures address various forms of data inconsistency—such as non-standardized station identifiers, missing geographic references, overlapping operational periods, and inconsistencies in connector type or CPO assignment. Particular attention is paid to detecting and resolving changes in station configuration over time, including geographic repositioning, rebranding, and infrastructural upgrades.

Then, we present the construction of two principal databases. The first is the *Infrastructure Database*, which contains detailed operational and spatial information for each charging station, including location, power level, plug type, operational timelines, and CPO affiliations. The second is a *Regional Aggregated Database*, which compiles region-level indicators such as the total number of charging stations, average power, vehicle fleet composition, and average EV traffic, thereby enabling spatial comparisons and policy-relevant insights.

To further understand the spatial distribution and demand potential of EV charging infrastructure across the country, the final section of this chapter introduces a quartile-based analysis of block traffic. This analysis aims to categorize charging stations according to the volume of vehicular flow in their immediate surroundings, thereby enabling a comparative assessment of station exposure and usage potential.

Block traffic is defined as the aggregated traffic flow within the hexagonal cells adjacent to the hexagon centered on the location of each charging station. This spatial structure allows for a more realistic estimation of the total traffic environment influencing each site, rather than limiting the measurement to a single point. These hexagonal units are part of a tessellated spatial grid used to discretize the study area, allowing for consistent spatial aggregation. For each charging station, we compute the sum of traffic values within the central hexagon and its immediate six neighbors, forming a first-order hexagonal ring. The traffic data were sourced from TomTom [52], a leading provider of high-resolution mobility and navigation data, and reflect annual traffic estimates across the Italian road network. By associating each station with its corresponding block traffic quartile, this analysis facilitates the identification of stations in high-demand areas and provides useful insights for infrastructure planning and investment prioritization.

This analysis is performed at two levels of territorial aggregation: first across the four Nielsen areas of Italy—North-West, North-East, Centre, and South—and subsequently across all individual regions. Stations are grouped into quartiles based on annual block traffic, defined as the number of vehicles passing in the vicinity of each station. This categorization offers a foundational understanding of spatial disparities in infrastructure exposure and utilization, serving as a key input for subsequent modeling efforts.

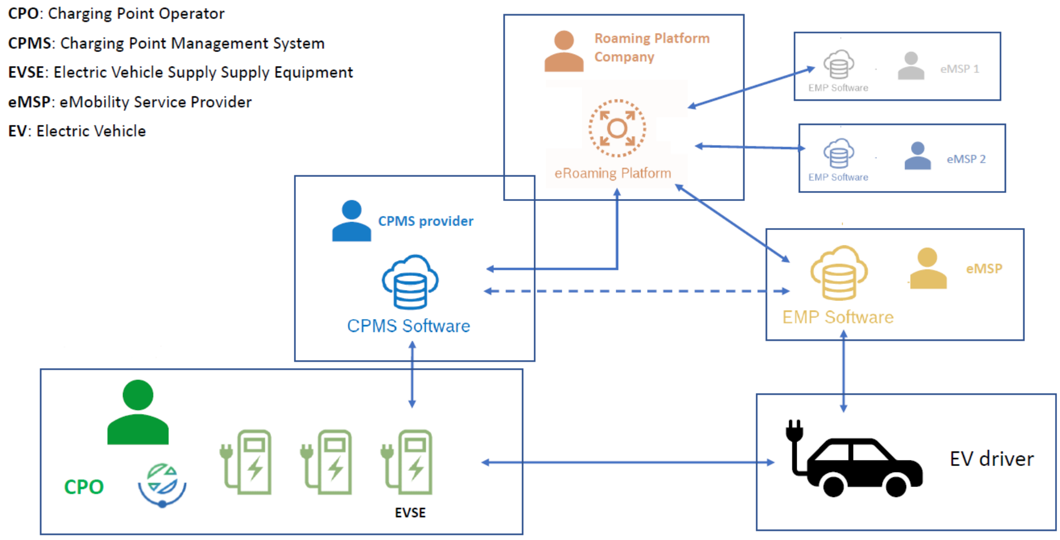

Figure 3.1: Main actors and information flows in the EV charging infrastructure.

3.1 Overview of EV Charging Infrastructure and Cost Structure

A proper understanding of electric vehicle (EV) charging infrastructure is essential before conducting data-driven analysis. This section introduces the basic elements involved in EV charging, including the physical components of the infrastructure (stations, connectors, and power levels), the technical standards regulating charging modes and plug types, and the key factors influencing electricity supply costs for charging sessions.

3.1.1 EV Charging Infrastructure Components

The electric vehicle (EV) charging ecosystem is composed of multiple actors and systems working together to ensure the availability, management, and accessibility of charging services. These components are fundamental not only for delivering electricity to EVs but also for enabling interoperability and real-time information exchange across networks. Figure 3.1 provides a schematic overview of the main actors involved.

The key entities in the EV charging infrastructure include:

- **CPO (Charging Point Operator):** The entity responsible for owning, installing, and maintaining EV charging infrastructure, also referred to as Electric Vehicle Supply Equipment (EVSE). CPOs manage the physical operation of charging stations and often collaborate with software providers to monitor charging sessions and station status.
- **CPMS (Charging Point Management System):** A software solution used by CPOs to control and monitor their EVSE assets. It collects data on station availability, session status (e.g., `OCCUPIED`, `AVAILABLE`), power delivery, and faults. The

CPMS interfaces directly with the charging hardware and acts as a bridge toward external systems like roaming platforms.

- **eMSP (eMobility Service Provider):** Companies that offer EV drivers access to a network of charging stations. They are responsible for customer authentication, billing, and providing apps or RFID cards for initiating charging sessions. eMSPs do not necessarily own charging infrastructure but provide access to it via agreements with CPOs or roaming platforms.
- **EV Driver:** The end-user who owns or operates an electric vehicle. EV drivers interact with the system through mobile apps or RFID cards provided by eMSPs to find, reserve, and initiate charging sessions.
- **Roaming Platform:** An intermediary infrastructure that facilitates data exchange and interoperability between multiple CPOs and eMSPs. It enables EV drivers to use charging stations operated by different CPOs without requiring multiple subscriptions, by relaying authentication, pricing, and session data in real time.
- **EVSE (Electric Vehicle Supply Equipment):** The physical charging station installed and operated by the CPO. It includes all hardware components such as connectors, meters, and communication modules.

This architecture allows seamless integration between hardware (EVSE), operational control (CPMS), customer-facing services (eMSP), and cross-network communication (roaming platforms). The distinction between these actors is particularly important when analyzing station usage patterns, pricing strategies, and data reliability. For instance, the data used in this thesis are collected via the partner's eRoaming platform, which aggregates real-time information from CPMS systems provided by CPOs across Europe.

3.1.2 Charging Standards and Connector Types

EV charging is governed by international standards that define both the hardware interface and the control systems used during charging. Two key IEC standards apply:

- **IEC 61851-1:** Defines general charging modes (AC and DC) and control protocols.
- **IEC 61851-23 / IEC 62196-1:** Define DC charging methods and physical connector types.

The most common plug types used in the dataset are described below:

- **Type 1:** A single-phase plug primarily used by American vehicles. It supports charging speeds up to 7.4 kW.
- **Type 2:** The European standard for EVs manufactured from 2018 onward. This triple-phase plug supports AC charging at speeds up to 43 kW.

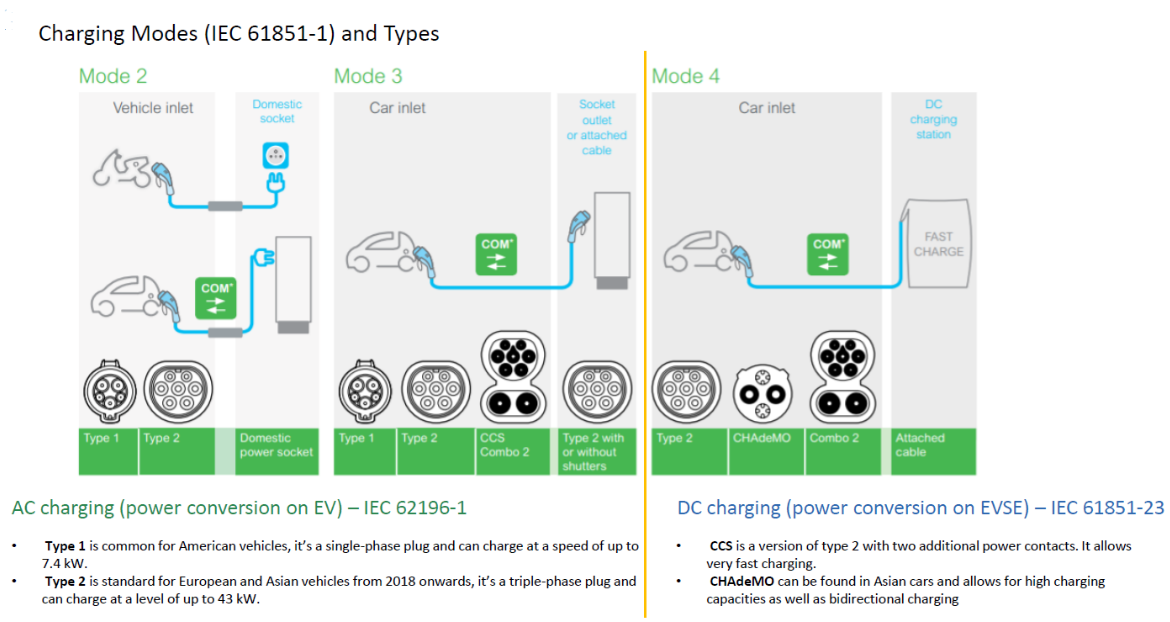

Figure 3.2: Summary of available connectors around the world.

- **CCS (Combined Charging System):** An extension of the Type 2 plug with two additional DC power contacts, enabling high-speed DC charging.
- **CHAdeMO:** A fast-charging DC standard commonly used in Asian vehicles. It supports high charging capacities and allows for bidirectional charging (vehicle-to-grid).

In AC charging, power conversion (from AC to DC) occurs inside the vehicle, while in DC charging, the power is converted within the charging infrastructure (EVSE), allowing for higher charging speeds.

3.1.3 Electricity Cost Structure for EV Charging in Southern Europe

The cost of providing electric vehicle (EV) charging services via public infrastructure depends on a complex combination of technical parameters, consumer classifications, and national regulatory frameworks. Grid tariffs, in particular, vary across countries and are influenced by factors such as voltage level, consumption characteristics, and whether the consumer is a Charge Point Operator (CPO) or a general non-domestic user.

Figure 3.3 summarizes the key variables influencing electricity costs in four Southern European countries—Italy, France, Spain, and Portugal—highlighting the differences in voltage-based pricing, time-dependency, and CPO-specific tariff schemes.

The structure of electricity pricing is typically defined by three core dimensions:

- **Voltage level:** Low Voltage (LV) or Medium Voltage (MV) connection affects eligibility for certain tariffs and power thresholds.

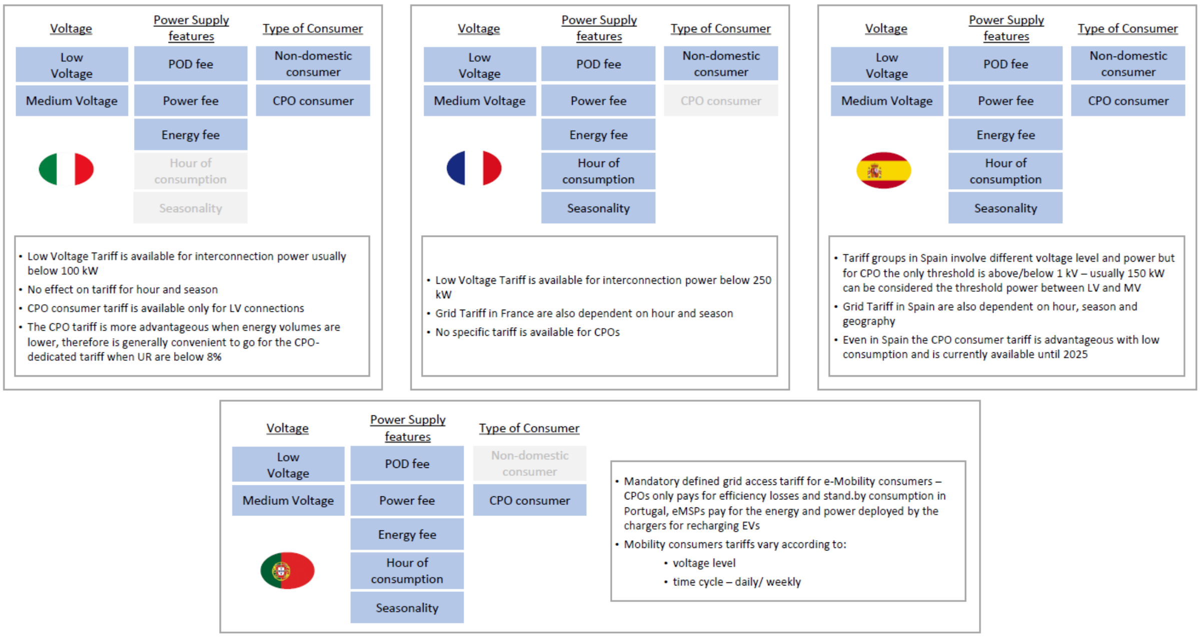

Figure 3.3: Summary of electricity cost components and CPO-specific tariff conditions across Italy, France, Spain, and Portugal.

- **Power supply features:** These include the POD (Point of Delivery) fee, power fee (based on contracted power), energy fee (based on actual consumption), and, in some cases, temporal variables such as hour of consumption or seasonal cycles.
- **Type of consumer:** Differentiates between general non-domestic consumers and CPOs, with the latter often eligible for dedicated pricing schemes.

Italy. In Italy, low-voltage tariffs apply to interconnection power levels typically below 100 kW, and no time-based (hourly or seasonal) differentiation is enforced. CPO-specific tariffs are only available under LV connections and tend to be more advantageous in cases of low energy utilization—particularly when the utilization rate (UR) is below 8%.

France. French grid tariffs allow LV connections for interconnection power below 250 kW and explicitly incorporate both hourly and seasonal components into the pricing scheme. However, no specific CPO-dedicated tariffs are currently available in France.

Spain. Spain differentiates between voltage levels and power thresholds, with a typical cut-off around 1 kV (approximately 150 kW). Grid tariffs are influenced by time-of-day, seasonal variation, and geographic location. A special tariff for CPOs exists and remains available until 2025; it offers economic advantages, particularly for low-consumption scenarios.

Portugal. Portugal has introduced a dedicated access tariff structure for e-Mobility consumers. CPOs are only charged for efficiency losses and standby consumption, while the energy and power consumed by users is billed to eMSPs (e-Mobility Service Providers). Tariffs vary by voltage level and temporal cycle (daily or weekly), offering a tailored pricing mechanism for EV charging infrastructure.

Understanding these tariff systems is essential for interpreting charging costs and assessing station profitability or accessibility in different regulatory environments. While this study focuses on data from Italy, the cross-country comparison provides valuable context for interpreting infrastructure deployment strategies and modeling real-world costs in broader European settings.

3.2 Leveraging eRoaming platform for Comprehensive Data

The eRoaming platform serves as a pivotal source of information for analyzing electric vehicle (EV) charging behavior. The platform aggregates detailed session-level data from its partner network of Charge Point Operators (CPOs), including unique connector identifiers, geospatial coordinates, operational statuses, and precise timestamps for session initiation and termination.

Nevertheless, it is essential to acknowledge the limitations inherent in the data collection process. As the eRoaming platform does not independently validate the information submitted by CPOs, the reliability of the dataset is contingent upon the accuracy and timeliness of third-party updates. Modifications reported by CPOs, such as alterations to charging point names, physical locations, or plug types, highlight the necessity for critical assessment of data fidelity when formulating operational or planning decisions.

This section provides a detailed overview of the original attributes present in the raw dataset, as received from the eRoaming platform, and the preprocessing and cleaning procedures applied to the data to ensure the data's suitability for downstream analysis, along with a description of the final dataset.

3.2.1 Data Extraction and Filtering Process

The dataset offers a comprehensive representation of Electric Vehicle Supply Equipment (EVSE) characteristics across a wide geographical span. Importantly, the analysis is conducted at the **connector level**, meaning that each record corresponds to a single charging connector rather than to the entire station. Key attributes include connector type, power rating, geographic coordinates, operational time frames, and CPO (Charge Point Operator) information, thereby providing a multidimensional view of the charging infrastructure landscape.

- **EVSE Identifier** (`evse_id`): the connector identifier, following a structured notation: the CPO country, its acronym, the station ID, and the connector ID.

- **Connector Status (`status`)**: the specific operative statuses:
  - **AVAILABLE**: The connector is ready for use.
  - **OCCUPIED**: The connector is currently in use.
  - **OUT_OF_SERVICE**: The connector is not operational.
  - **RESERVED**: The connector is reserved for a specific user.
  - **UNKNOWN**: The status of the connector is not known.
- **Status Duration (`status_duration_in_seconds`)**: measures the time in seconds that the connector has been in its current status.
- **Plug Type (`plugs`)**: the types of plug of the connector. The types include: CCS Combo 2 Plug (Cable Attached), CHAdeMO, Type 3 Outlet, Type 2 Outlet.
- **Geographical Information**: the dataset includes the city name, latitude (`lat`), and longitude (`long`) for each connector, aiding in geographical analysis and mapping.
- **Date of Record (`aggregated_date`)**: timestamping each record.

The initial step of the analysis involves isolating sessions associated with DC (Direct Current) charging stations. This is accomplished by filtering the dataset based on specific plug types that are indicative of DC fast charging, namely, the **CCS Combo 2 Plug (Cable Attached)** and **CHAdeMO**. This focus aligns with the strategic direction of the company Atlante, which has prioritized the deployment of DC charging infrastructure as a core element of its business model.

Following the identification of DC charging sessions, the dataset is further refined to retain only those records corresponding to the **OCCUPIED** status to filter charging sessions actively in progress. To improve the accuracy of session reconstruction, consecutive "OCCUPIED" statuses are merged into continuous charging events, even when they are interrupted by short intervals of up to 30 seconds, marked as "UNKNOWN" or "OUT_OF_SERVICE". Such brief interruptions are assumed to reflect temporary system or connector errors rather than genuine session terminations.

3.2.2 Data Validation, Cleaning and Mapping Methodology

Preparing the charging station dataset for analysis required a structured three-phase workflow: validation and filtering, cleaning, and mapping. Filtering focuses on selecting relevant records based on charging type and session status. Cleaning ensures consistency and accuracy across identifiers, geospatial data, and operational timelines. Finally, mapping addresses the identification of structural and temporal changes, enabling the reconstruction of station evolution over time.

Data Validation

To enhance the accuracy and reliability of the charging point data, a comprehensive multi-step validation procedure was adopted. This methodology systematically verifies the integrity of information recorded on the eRoaming platform, particularly in relation to reported modifications. The validation process comprises several key components:

1. **Verification of Change Duration**: assessing the temporal validity of each reported change to a charging station's attributes. Only those modifications that persist for a minimum of 30 days beyond the commissioning date are retained. This criterion filters out transient or erroneous updates.

2. **Assessment of Change Sequentiality**: to preserve the chronological consistency of the dataset, the validation process includes a check on the sequential order of recorded changes. Specifically, the end date of one configuration must precede the start date of the subsequent one. This requirement prevents overlapping timeframes and maintains the temporal integrity of each charging point's data history.

3. **Validation of CPO Transition Consistency**: in cases where a charging station changes management, such as the transition from Enel X to Ewiva, a joint venture with Volkswagen, additional scrutiny is applied. When the geographical coordinates remain constant, the change in Charge Point Operator (CPO) is verified for logical consistency.

Data Cleaning

The cleaning of Electric Vehicle (EV) charging station data represents a crucial step in ensuring the integrity, consistency, and analytical value of the dataset. This process aims to refine connector-level information by addressing issues related to geospatial precision, connector type standardization, and the temporal consistency of operational status records.

The cleaning phase consists of the following steps:

1. **Standardization of EVSE IDs:** Unique identifiers for connectors (EVSE IDs) are normalized to conform to a standardized format in case they did not follow the standard notation.

2. **Geospatial Data Validation:** Location data, such as geographical coordinates and city names, are validated and corrected using an official shapefile of Italy. This process guarantees that each station is correctly positioned within its administrative boundaries, thereby supporting reliable spatial analysis.

3. **Grouping by Unique Attributes:** connectors are grouped according to key identifiers, including EVSE ID, connector types, and geographic coordinates. This aggregation is instrumental in detecting and removing duplicate or fragmented entries.

4. **Temporal Validation:** Changes in connector attributes over time are examined for logical and chronological consistency. This involves verifying the duration and sequence of each modification to ensure that the data reflects a coherent timeline of station updates.

Mapping and Identification of Changes

A key component of the methodology focuses on identifying temporal and structural changes in charging station records by mapping and comparing entries within the cleaned dataset. This phase aims to detect inconsistencies and track the evolution of infrastructure elements, particularly where changes are subtle or partially documented. The mapping process involves the following procedures:

1. **Construction of a Mapping DataFrame:** A dedicated mapping DataFrame was developed to trace modifications in spatial coordinates, connector specifications, and other operational attributes. This structured representation facilitates systematic comparisons and enables the identification of changes across time for individual connectors.

2. **Detection of Repositioned Connector:** Instances where the same `EVSE_ID` appears with differing latitude and longitude values were examined. In such cases, we verified whether the associated connectors remained consistent by calculating inter-record distances and evaluating the continuity of operational dates. This step helps distinguish between genuine IDs' reassignment and erroneous coordinate updates.

3. **Identification of Duplicate Connectors:** Records with nearly identical geographic coordinates but different `EVSE_IDs` were flagged as potential duplicates or reassignments. This was particularly important for detecting instances where a charging point may have been re-registered under a new identifier without a significant physical relocation.

4. **Plug Name Variation Analysis:** Charging stations with the same `EVSE_ID`, latitude, and longitude but differing plug names were also identified. This analysis captures updates in connector types or reclassifications, which are critical for understanding infrastructure evolution and compatibility shifts.

3.2.3 eRoaming Platform Infrastructure Database

The finalized eRoaming platform database represents a comprehensive and structured compilation of cleaned and validated data on individual connectors. A detailed overview of the main fields included in the database is provided below:

- **EVSE ID (`evse_id`):** A standardized unique identifier for each charging station, typically structured as a combination of country code, CPO acronym, station ID, and connector ID.

- **Plugs (plugs):** Specifies the type(s) of connectors available at the station.
- **Geographical Information:** Includes latitude (lat), longitude (lon), and city name (city).
- **Operational Time Frame:** Captures the first and last observed dates of activity (first_date and last_date), delineating the station's operational lifespan.
- **Connector Type (connettore):** Indicates the charging standard (e.g., Type 2, CCS) used at the station.
- **Charger and Operator Metadata:** Includes internal charging point identifiers (evse_id_charger) and a standardized name for the Charge Point Operator (cpo_clean).
- **Administrative Information:** Identifies the administrative region (den_reg), province (den_prov), and municipality (comune) in which the station is located.
- **Additional Attributes:** Encompasses supplementary variables such as the ecomovement identifier (evse_id_eco), various proximity metrics (proximity, station_n_prox_cpo, station_cl), and a concatenated station label (station_name).

Example Entry

An illustrative example from the database is presented to elucidate the structure and type of data recorded:

```
evse_id: IT*ELX*E21BZ0807B*1
plugs: ["CCS Combo 2 Plug (Cable Attached)"]
city: San Vitaliano
lat: 40.9358
lon: 14.4898
first_date: 2022-01-01 14:14:00.000
last_date: 2022-01-02 15:18:00.000
connettore: 1
evse_id_charger: IT*ELX*E21BZ0807B
den_reg: Campania
den_prov: -
comune: San Vitaliano
evse_id_eco: ENELX
cpo_clean: ENELX
lat_rad: 0.7144644919378947
lon_rad: 0.25289471795547436
```

```
proximity: 0
station_n_prox_cpo: 0
station_cl: ENELX_San Vitaliano_0
```

3.2.4 Database Splitting: Dimensions and Facts

After thorough consideration, the comprehensive connector-level database was first aggregated into a station-level dataset to improve interpretability and facilitate the subsequent modeling efforts. This process began by retrieving the power level associated with each connector using its unique identifier. Connectors were then grouped by their corresponding charging station ID and power level, allowing us to compute the number of connectors available at each power level for a given station. As a result, each entry in the final dataset is uniquely defined by a combination of station ID and connector power. For clarity and simplicity, we will refer to these aggregated entities as "stations" throughout the remainder of the thesis, even though they technically represent station-power level pairs.

Following this aggregation, the dataset was logically bifurcated into two distinct segments to enhance data organization and analytical flexibility: a *dimensional database*, containing static attributes such as geographic location, administrative boundaries, and CPO information; and a *factual database*, which includes time-dependent variables such as session counts, operational statuses, and calculated usage metrics. This separation between dimensions and facts enables more efficient data querying and supports both descriptive and predictive analyses.

Dimension Database

The Dimension Database stores attributes that are invariant over time. These include both original data fields and additional derived features. Using the basic available information for each charging station, namely the geographic coordinates (latitude and longitude), commissioning and decommissioning dates, and unique station identifiers, we were able to enrich the dataset with a series of informative attributes that provide greater spatial and functional context. These include:

- **on_the_go:** A binary indicator denoting whether the charging station is located in a context typically associated with transient charging behavior (e.g., roadside or en-route locations).
- **operation_days:** Total number of days the station was active, calculated from the commissioning and decommissioning timestamps.
- **distance_to_primary:** Distance (in kilometers) from the station to the nearest primary road, providing insight into accessibility and visibility.
- **motorway_flag:** A binary indicator denoting whether the station is located near a motorway, which may affect charging behavior and usage profiles.

- **count_in_site_DC:** Number of DC connectors available in a radius of 500 meters, providing a local density metric for fast-charging availability.

These additional features complement the static station information already included in the database and are designed to improve both descriptive analysis and model performance in subsequent stages, which are listed below:

- **station_cl:** Unique identifier for the charging station obtained with the complete CPO name and information retrieved through the coordinates.
- **conn_power:** Power output capability of the charging station in kW.
- **station_id_power:** Composite identifier combining the name of the station in the station_cl column and its power output.
- **n_cp:** Number of charging points available at the station.
- **cpo_clean:** Name of the Charge Point Operator (CPO) after data cleaning procedures.
- **single_traffic** and **block_traffic:** Traffic metrics around the charging station, indicating the volume of potential EV drivers.
- **perc_rank_single_traffic** and **perc_rank_block_traffic:** Percentile rank of the station based on traffic metrics, providing insights into the station's accessibility and potential demand.
- **population_prov, population_reg, population_comune:** Population metrics within the province, region, and municipality, respectively, of the charging station.
- **commissioning** and **de_commissioning:** Dates indicating when the station was commissioned (first charging station registered) and decommissioned (last charging station registered), offering a lifespan view of the station.
- **den_reg** and **comune:** Administrative region and municipality where the station is located.
- **lat** and **lon:** Geographic coordinates of the charging station.

Example Entry

```
station_cl: A2A_San Lorenzo di Sebato_1432
conn_power: 63.0
station_id_power: A2A_San Lorenzo di Sebato_1432_63.0
n_cp: 1
```

```
cpo_clean: A2A
single_traffic: 4760.0
perc_rank_single_traffic: 0.18188786684187472
block_traffic: 17100.0
perc_rank_block_traffic: 0.26872536136662284
population_prov: 656924
population_reg: 1389815
population_comune: 4316
commissioning: 2022-04-23 08:07:00.000
de_commissioning: 2023-11-29 20:01:00.000
den_reg: Trentino-Alto Adige
comune: San Lorenzo di Sebato
lat: 46.7754
lon: 11.8799
```

Facts Database

The Facts Database contains attributes that vary over time, specifically the information about the charging sessions of the charging stations:

- **station_cl, conn_power:** Same as the Dimensional Database columns.
- **year, month, day:** Temporal markers indicating the specific date of the charging session.
- **station_id_power:** Same as the Dimensional Database column.
- **min_days_in_month:** The minimum number of days in the month being reported, useful for calculating average daily metrics.
- **total_sessions:** Total number of charging sessions recorded on the specified day.
- **avg_session_over_connectors:** Average number of sessions per connector, providing insights into utilization rates.
- **total_or_daily** and **avg_or_daily_over_connectors:** Metrics related to the Occupancy Rate (OR), both in total and averaged over the number of connectors, indicating the station's usage intensity.
- **distinct_evse_ids:** The number of unique EVSE IDs (so charging point) that the charging station has.

3.3 EV Infrastructure and Adoption Metrics

3.3.1 Regional Data Overview

Leveraging the newly constructed dataset, a regional-level dataset was developed to support preliminary analyses of the electric vehicle (EV) market across Italy. The analysis is based on data from the year 2023, in alignment with the temporal scope of the eRoaming platform dataset, to ensure coherence across all data sources. This regional dataset integrates several key indicators that reflect both the current state of EV infrastructure and its potential for future expansion.

In particular, data on the number of total and electric vehicles registered in each region were retrieved from the open data portal of the Automobile Club d'Italia (ACI) [6], providing an official and reliable foundation for assessing EV adoption levels. The integration of these vehicle statistics with infrastructure data enables a more comprehensive evaluation of supply–demand dynamics in different parts of the country. The variables included in the regional dataset are as follows:

- **Total Stations**: Represents the absolute count of charging points within each region.
- **% Electric Stations**: Indicates the share of electric charging infrastructure relative to total charging facilities.
- **Average Power Output**: Provides a regional mean of the power capacity (in kW) delivered by available charging stations.
- **Vehicle Stock** (`Total Cars` and `Electric Cars`): Captures the total number of registered vehicles and the subset that are electric, respectively.
- **% Electric Vehicle over Total**: Denotes the percentage of electric vehicles within the overall vehicle fleet in each region.
- **Average Daily EV Traffic** (`Single Traffic Average` and `Block Traffic Average`): Measures the average daily flow of electric vehicles accessing charging infrastructure, both individually and in traffic blocks.

This regionalized view provides valuable insights into the spatial heterogeneity of EV infrastructure and usage patterns across Italy, supporting both strategic planning and policy evaluation.

Infrastructure Distribution and Electrification Rate. The total number of charging stations varies markedly among regions, with *Lombardia* (729 stations, 22.42% of the total) and *Lazio* (365 stations) emerging as the most infrastructurally developed areas. These two regions alone account for a substantial share of the national public charging network, reflecting their high population density and economic relevance. Conversely, regions such as *Molise* and *Basilicata* report very limited coverage (8 and 11 stations, respectively), suggesting clear territorial disparities in accessibility to EV infrastructure. This imbalance highlights the geographical concentration of electrification efforts in northern and central Italy, with southern regions lagging behind in the diffusion of charging infrastructure.

Power Capacity. The average power output per station ranges between 50 kW (*Molise*) and nearly 100 kW (*Abruzzo*), indicating moderate heterogeneity in the technological level of the infrastructure. Regions with higher average power values, such as *Abruzzo* and *Umbria*, likely feature a larger share of fast or ultra-fast charging points, improving service efficiency and user experience. In contrast, regions such as *Sardegna* and *Sicilia* present lower power averages (around 60–70 kW), which could result in longer charging times and reduced station attractiveness.

Vehicle Stock and Electrification Penetration. The number of registered vehicles also exhibits strong regional asymmetry. *Trentino-Alto Adige* (2.15%) and *Valle d'Aosta* (1.16%) stand out, displaying a relatively high level of EV penetration despite their limited vehicle fleets. This pattern could reflect the effectiveness of local policies and incentives that have encouraged early adoption. In contrast, regions such as *Calabria*, *Puglia*, and *Basilicata* remain below 0.15%, underscoring the persistent north–south divide in the Italian EV market.

Traffic and Utilization Patterns. Traffic-related indicators provide further insight into the spatial dynamics of charging demand. *Trentino-Alto Adige*, *Lazio*, and *Lombardia* exhibit the highest average daily EV traffic, both in individual and block measures, with block averages exceeding 380 in *Lombardia* and 630 in *Trentino-Alto Adige*. These results suggest not only greater utilization of the charging network but also higher integration of EVs within regional mobility systems. Conversely, *Molise* and *Basilicata* display minimal traffic activity, consistent with their limited infrastructure and low EV stock.

Table 3.1: Regional Data Overview

Region	Total Stations	% Electric Stations	Avg. Power	Total Cars	Electric Cars	% Electric Cars over Tot	% Electric Cars EV	Single Traffic Avg.	Block Traffic Avg.
Abruzzo	65	2.0	99.74	903081	2012	0.22	1.27	73.91	127.69
Basilicata	11	0.34	87.5	383305	486	0.13	0.31	9.82	16.27
Calabria	34	1.05	72.22	1338121	1366	0.10	0.86	34.20	65.03
Campania	107	3.29	94.5	3612878	4438	0.12	2.81	46.58	89.43
Emilia-Romagna	305	9.38	93.14	2961375	10707	0.36	6.77	88.12	191.08
Friuli Venezia Giulia	48	1.48	85.425	812503	2798	0.34	1.77	34.58	58.06
Lazio	365	11.22	74.68	3857390	17604	0.46	11.13	173.74	386.93
Liguria	93	2.86	88.56	843142	2166	0.26	1.37	76.00	134.83
Lombardia	729	22.42	73.86	6272187	31429	0.5	19.88	154.04	381.95
Marche	107	3.29	88.59	1043160	2894	0.28	1.83	79.34	167.06
Molise	8	0.25	50.0	215043	279	0.13	0.18	6.375	21.37
Piemonte	333	10.24	76.45	2900449	10360	0.36	6.55	78.88	184.90
Puglia	91	2.80	96.86	2451311	3212	0.13	2.03	43.23	87.20
Sardegna	96	2.95	61.06	1097782	2309	0.21	1.46	56.0	108.51
Sicilia	139	4.27	66.85	3438078	4795	0.14	3.03	20.85	43.00
Toscana	181	5.57	79.71	2634922	15373	0.58	9.72	144.34	300.51
Trentino-Alto Adige	228	7.01	78.47	1276378	27420	2.15	17.34	338.93	630.58
Umbria	48	1.48	98.08	646307	1567	0.24	0.99	105.79	138.83
Valle d'Aosta	38	1.17	80.69	287951	3331	1.16	2.11	87.10	132.05
Veneto	226	6.95	88.89	3221693	13585	0.42	8.59	94.79	193.58

3.3.2 Nielsen Areas Data Overview

Table 3.2: Summary of charging infrastructure by Nielsen area

Area	Total stations	% Stations over Total	Avg. Pwr (kW)	Block traffic
Centre	701	21.6	79.8	314
North-East	807	24.8	87.2	308
North-West	1193	36.7	76.1	300
South	551	16.9	80.1	81

Aggregate distribution across Nielsen areas. Table 3.2 shows that the North-West leads in charging infrastructure with 1193 total charging stations (36.7% of the national stock in Nielsen areas), followed by the North-East with 807 charging stations (24.8%), the Centre with 701 charging stations (21.6%), and the South (Sud) trailing at 551 charging stations (16.9%).

Charging power and fleet electrification. Average station power varies from 76.1 kW in North-West to 87.2 kW in North-East, with Centre (79.8 kW) and Sud (80.1 kW) in between. Although North-West hosts the largest number of stations, its average power is the lowest, suggesting many medium-power nodes. In terms of vehicle stock, Sud has by far the largest fleet (13.44 million cars) but the lowest EV penetration (0.14% of total cars), compared with North-East (0.66%), Centre (0.46%) and North-West (0.46%). This indicates that despite a high overall fleet size, the South remains under-electrified relative to the North.

Usage intensity and session patterns. Looking at average daily charging sessions, North-East records the highest usage per column with 158 single sessions and 308 block sessions per day, closely followed by Centre (147/314) and North-West (125/300). South lags behind with just 42 single and 81 block sessions. The ratio of block to single sessions is highest in Centre ($314/147 \approx 2.1$), hinting at more clustered charging events, whereas South's ratio ($81/42 \approx 1.9$) remains modest despite low absolute volumes.

3.4 Analysis of Quartile Distribution of EV Charging Stations in Italy

3.4.1 Regional analysis

The data outlines the distribution of electric vehicle (EV) charging stations across various regions in Italy, categorized into three quartiles based on block traffic, which signifies the number of cars passing nearby the station annually.

Table 3.3: Percentage distribution of charging stations by annual block-traffic quartiles per region

Region	Low traffic (%)	Medium traffic (%)	High traffic (%)
Lombardia	27.89	30.03	42.08
Piemonte	32.15	32.01	35.84
Trentino-Alto Adige	62.43	21.89	15.68
Lazio	24.53	29.69	45.77
Emilia-Romagna	27.98	40.21	31.81
Sicilia	62.31	23.13	14.55
Veneto	34.84	41.40	23.76
Toscana	25.92	37.90	36.19
Valle d'Aosta	97.40	2.60	0.00
Friuli-Venezia Giulia	72.63	27.37	0.00
Marche	30.37	33.18	36.45
Sardegna	40.85	35.92	23.24
Campania	16.54	40.77	42.69
Liguria	17.50	52.50	30.00
Puglia	16.38	48.02	35.59
Molise	100.00	0.00	0.00
Basilicata	100.00	0.00	0.00
Abruzzo	15.91	44.32	39.77
Calabria	16.00	48.00	36.00
Umbria	7.41	54.63	37.96

Overall trends. The results confirm a marked territorial heterogeneity in the utilization of charging infrastructure. Northern and central regions—particularly *Lazio*, *Lombardia*, and *Emilia-Romagna*—display a balanced or high concentration of stations within the upper traffic quartiles, whereas southern and insular regions are characterized by a predominance of low-traffic sites.

High-traffic regions. *Lazio* (45.77%), *Campania* (42.69%), and *Lombardia* (42.08%) show the largest shares of high-traffic stations. These values are consistent with their elevated levels of electric-vehicle penetration and overall network density, suggesting the existence of agglomeration effects where infrastructure availability and user demand reinforce each other. In *Toscana* and *Emilia-Romagna*, high-traffic stations also represent around one third of the total, indicating an emerging but less polarized demand pattern.

Medium-traffic regions. Regions such as *Umbria* (54.63%), *Liguria* (52.50%), and *Puglia* (48.02%) register the majority of their infrastructure in the medium-traffic quartile. This configuration implies a moderate yet widespread level of utilization, possibly reflecting networks that are neither saturated nor underused. The distribution suggests stable regional adoption dynamics that could intensify with further policy or market incentives.

Low-traffic regions. At the opposite end, *Trentino-Alto Adige* (62.43%), *Sicilia* (62.31%), and *Friuli-Venezia Giulia* (72.63%) exhibit a predominance of low-traffic stations. Despite the strong policy commitment in *Trentino-Alto Adige*, the high share of low-traffic chargers may be due to the dispersion of infrastructure in peripheral or touristic areas with intermittent demand. *Molise* and *Basilicata* show exclusively low-traffic networks (100%), confirming their marginal role within the national EV ecosystem.

Interpretation. The data reveal a clear core–periphery pattern: metropolitan and industrialized regions host a significant proportion of high-traffic stations, while peripheral and southern territories remain dominated by low-usage sites. This spatial imbalance highlights the necessity of differentiated planning strategies. In mature markets (e.g., *Lombardia*, *Lazio*), the focus should shift toward upgrading power capacity and reducing congestion, whereas in emerging areas, priority should be given to demand stimulation through incentives and strategic siting. Overall, the results demonstrate that utilization intensity is not merely a function of infrastructure quantity but also of locational and socio-economic context, reinforcing the need for spatially informed EV infrastructure policies in Italy.

3.4.2 Nielsen Areas

Table 3.4 reports the distribution of charging stations across annual block-traffic quartiles, aggregated by Nielsen area. This representation offers a macro-regional perspective on the spatial patterns of charging infrastructure utilization in Italy.

Table 3.4: Distribution of charging stations by annual block-traffic quartiles and Nielsen area

Area	% Low traffic	% Medium traffic	% High traffic
North-West	30.69	31.71	37.60
North-East	40.15	35.77	24.08
Centre	25.85	34.15	40.00
South	34.17	35.80	30.03

Overall trends. The results highlight significant regional disparities in station usage intensity. The *Centre* and *North-West* exhibit the highest shares of high-traffic stations (40.00% and 37.60%, respectively), while the *North-East* shows a markedly lower proportion (24.08%) and a predominance of low-traffic sites (40.15%). The *South* presents a more balanced distribution, with roughly one third of stations in each traffic class.

Interpretation by macro-area. The prevalence of high-traffic stations in the *Centre* and *North-West* suggests stronger EV adoption and greater infrastructure maturity, likely reflecting higher population density, economic activity, and the presence of metropolitan areas such as Rome, Milan, and Turin. These areas combine dense charging networks with sustained demand, indicating that infrastructure investment and vehicle uptake are evolving in tandem. In contrast, the *North-East*, despite being industrially developed, displays a concentration of low-traffic stations. This could be linked to a more dispersed spatial distribution of chargers and a comparatively lower share of urban users. The *South* demonstrates a nearly uniform allocation across traffic categories, reflecting a transitional phase in EV adoption. While the share of high-traffic stations (30.03%) is not negligible, the relatively balanced distribution indicates that utilization is still heterogeneous and context-dependent, influenced by local mobility patterns and the limited density of charging infrastructure.

3.5 Points of interest data: two approaches

In addition to the infrastructure and usage data obtained from the eRoaming platform, this study incorporates a rich layer of contextual information derived from OpenStreetMap (OSM) [39]. These external data sources provide essential insights into the urban and functional environments surrounding each charging station—factors that are known to influence user behavior, station attractiveness, and overall infrastructure effectiveness.

3.5.1 Deviation Flow Model Approach

To capture these spatial dynamics, the contextual data were integrated into two distinct modeling frameworks, each requiring a tailored representation of the built environment. The first framework is based on the principles of the *Deviation Flow Model*, which is commonly used in location analysis to account for the possibility that individuals may deviate from their intended path to access a service. While the behavioral drivers underlying such deviations are often difficult to quantify—particularly in the case of EV drivers, whose decision-making may be shaped by vehicle range, urgency, or convenience—this model provides a structured way to incorporate the potential influence of nearby amenities on station usage. The assumption, supported in part by analogous studies in retail and transportation systems, is that a sufficiently attractive set of surrounding services (e.g., shops, restaurants, or recreational facilities) may justify a minor detour from the optimal route.

In this context, Points of Interest (POIs) were grouped into thematic categories—such as *leisure*, *hospitality*, *institutional*, *shopping*, and *logistics*—and quantified within 5- and 10-minute catchment areas, by both foot and car. This proximity-based feature engineering allows us to approximate the accessibility and diversity of services available around each station. Notably, malls were treated as a standalone category, as an initial manual inspection of the highest-performing stations revealed a strong spatial correlation with large commercial centers, suggesting that such destinations may exert a unique pull on EV users.

3.5.2 Node Serving Problem Approach

The second approach addresses the *Node Serving Problem*, which reframes the analysis from a route-level to a network-level perspective. In this phase, the dataset was constructed at the node level, where each node represents a unique spatial entity—such as a city or municipality—eliminating redundancy and avoiding the overlap inherent in proximity-based buffers. A more granular and mutually exclusive POI classification was adopted, comprising categories such as *supermarket*, *healthcare*, *fast_food*, *office*, *tourism*, and others. This finer-grained classification supports a detailed characterization of urban functionality and land use around each node, enabling more accurate modeling of the factors that influence whether a node is adequately served by the existing charging infrastructure. The categories considered include:

- **Retail and Commercial Spaces**: Locations where users might stop for shopping, dining, or other commercial activities. This includes shops, restaurants, cafes, and convenience stores. Certain large-scale commercial areas, such as malls and department stores, were excluded to avoid overrepresentation of high-traffic locations.
- **Residential and Accommodation Facilities**: Housing units, apartment complexes, and hotels that can serve as potential sources of demand, particularly for overnight or long-duration charging.
- **Recreational and Leisure Locations**: Parks, entertainment centers, and cultural institutions such as museums, theaters, and sports venues that attract visitors and may create short-term demand for charging.
- **Tourism-Oriented POIs**: Hotels, hostels, and tourist attractions such as historical landmarks, galleries, and theme parks that could lead to increased charging needs, particularly in areas with significant tourist footfall.
- **Transport and Mobility Hubs**: Locations such as train stations, airports, and major bus terminals that facilitate intermodal transport, potentially increasing the need for EV charging among commuters and travelers.
- **Healthcare Facilities**: Hospitals, clinics, and pharmacies that could influence the need for EV charging, particularly for staff and visitors.
- **Public Services and Institutions**: Government buildings, post offices, libraries, and educational institutions that attract daily commuters and long-duration visitors.
- **Financial Services**: Banks, ATMs, and other financial institutions that may contribute to short-duration stops.

Table 3.5: POI Categories Used in the Deviation Problem Analysis

Category	Description
Leisure	Includes amenities such as parks, sports facilities, and recreational venues that users may be willing to visit during a charging session.
Institutional	Encompasses public services and administrative buildings, such as schools, government offices, and libraries.
Hospitality	Covers POIs such as hotels, cafes, and restaurants, which are particularly relevant for longer charging sessions.
Shopping	Refers to general retail locations excluding malls, including grocery stores, boutiques, and standalone retail shops.
Mall	Kept as a separate category due to its distinctive role identified in preliminary analyses. Manual inspection of high-performing charging stations revealed a strong spatial association with large shopping malls, suggesting they serve as major attractors for EV users.
Transportation	Includes POIs related to mobility, such as bus stops, train stations, and taxi stands.
Logistics	Involves distribution centers, warehouses, and service depots, often relevant in industrial or high-traffic zones.

- **Spiritual and Religious Sites**: Churches, mosques, temples, and other places of worship that may generate periodic demand for charging infrastructure.

Methodology for Data Retrieval

To systematically assess the presence and density of POIs near charging stations, the following approach was adopted:

1. **Definition of Search Area**: A radius of 300 meters was chosen to capture POIs in proximity to each charging station, ensuring that the dataset reflects realistic stopping behavior.

2. **Data Extraction from OpenStreetMap (OSM)**: The OSM database was used as the primary source due to its comprehensive global coverage and detailed categorization of POIs. The retrieval process focused on extracting features tagged with relevant attributes, such as `shop`, `building`, `leisure`, `tourism`, and `transport`.

3. **Categorization and Filtering**: Extracted POIs were categorized into the predefined groups. Filtering criteria were applied to remove redundant or irrelevant features, ensuring a meaningful dataset.

4. **Spatial Analysis**: Each charging station's surroundings were analyzed to determine the density and variety of POIs. This analysis helps identify potential relationships between POI distribution and station utilization rates.

Data Processing and Interpretation

The retrieved dataset underwent further processing to enhance its usability and accuracy:

- **Normalization of POI Categories**: Ensuring consistency across diverse data sources by standardizing POI classifications.
- **Aggregation by Type and Density**: Grouping POIs by category and calculating their density within the defined radius for comparative analysis.
- **Exclusion of Irrelevant Features**: Filtering out large-scale commercial centers, duplicate entries, and misplaced data points that could distort results.

This structured approach ensures that the dataset accurately reflects the interaction between POIs and EV charging behavior, enabling further analysis on how different types of locations influence charging station utilization.

Final Features in the Dataset

The final dataset includes the following key features, each representing different POI categories:

Table 3.6: POI Categories Used in the Node Serving Problem

Category	Description
shop	General retail shops excluding specialized categories (e.g., clothing stores, electronics).
mall	Large shopping complexes often encompassing multiple retail and food outlets.
department_store	Multi-category retail establishments that are distinct from malls in layout and scale.
supermarket	Grocery-focused retail outlets offering daily-use products.
wholesale	Bulk-purchase commercial centers typically catering to businesses or high-volume customers.
furniture	Retail locations specializing in home furnishings and decor.
sustenance_short	Quick-service food options such as bakeries, cafes, or takeaway-only outlets.
sustenance_long	Full-service restaurants or dining establishments requiring longer stays.
education	Schools, universities, training centers, and other academic institutions.

Continued on next page

Category	Description
transportation	Facilities such as bus stops, train stations, airports, and taxi stands.
financial	Banks, ATMs, insurance offices, and other financial service points.
healthcare	Hospitals, clinics, pharmacies, and medical service providers.
entertainment	Cinemas, theaters, concert venues, and other recreational locations.
public_service	Government buildings, municipal offices, and administrative service points.
amenity_facilities	Utility-based infrastructure such as public toilets, drinking water stations, and charging points (non-EV).
amenity_others	Other amenities not classified under basic facilities, such as community centers or event spaces.
spiritual	Religious sites including churches, mosques, synagogues, and temples.
fast_food	Chain-based quick-service restaurants with standardized menus and minimal seating time.
fuel	Traditional fuel stations (e.g., petrol, diesel), relevant for evaluating mixed-infrastructure environments.
parking	Designated parking facilities including surface lots, garages, and park-and-ride areas.
house	Residential POIs including standalone homes, housing complexes, or residential towers.
sports	Sports facilities such as gyms, stadiums, swimming pools, and courts.
office	Business and administrative office spaces.
leisure	Parks, green areas, playgrounds, and recreational open spaces.
tourism	Tourist attractions such as monuments, museums, scenic viewpoints, and cultural heritage sites.

3.6 Construction of a Synthetic Spatio-Temporal Demand for EV Charging

In the absence of complete and high-resolution traffic or charging-demand data, it becomes necessary to approximate mobility patterns through a combination of demographic information, theoretical models, and empirical observations. The objective of this section is to construct a realistic spatio-temporal representation of electric vehicle (EV) charging demand across the study area, suitable for downstream optimisation and planning analyses. To achieve this, we integrate multiple data sources and modelling steps: (i) high-resolution European population grids, used to estimate the potential user base around each network node; (ii) the Radiation Model, which provides a theoretically grounded mechanism to translate population distributions into plausible origin–destination flows; (iii) observed mode shares and temporal mobility profiles derived from the Regione Lombardia OD matrix, used to calibrate the flows and incorporate realistic time-of-day variations; and (iv) empirical charging-session distributions by connector power class obtained from Atlante, enabling the final allocation of demand across different charging technologies. By combining these elements, we obtain a coherent, data-driven and temporally resolved demand model that reflects both theoretical mobility principles and observed behavioural patterns, despite the lack of direct traffic measurements.

3.6.1 Population Estimation Around Graph Nodes

As the first step, we constructed a realistic proxy by estimating the population potentially interacting with each node of the network. This was achieved by integrating high-resolution population grid data with the spatial distribution of the nodes. The approach combines spatial buffering, geometric intersections, and area-weighted aggregation to compute a population measure consistently associated with each point in the graph.

The population grid data were obtained from the GEOSTAT dataset provided by Eurostat [18], where each cell is encoded in `EPSG:3035` coordinates and contains the total resident population in 2021. Each grid identifier embeds essential geometric information, such as the cell resolution and the coordinates of its centroid. Based on this information, we reconstructed the geometry of each cell by generating the corresponding square polygon.

After reconstructing the population grid, the data were transformed into a metric coordinate system to enable spatial operations. Each node of the network was then represented as a point and expanded into a circular buffer of fixed radius (1 km). This buffer represents the local catchment area of the node, i.e., the region from which users could realistically contribute to charging demand or mobility flows.

The buffered nodes were intersected with the population grid to determine which portions of the population cells lie within each catchment area. Because buffers and grid cells do not overlap perfectly, each intersection polygon captures only a fraction of the cell area; therefore, we computed a proportional share of the cell's population based on the ratio between the overlapping area and the total area of the grid cell. Summing these weighted contributions across all intersecting cells yields an estimate of the population influenced by, or served by, each node.

The resulting variable, `population_buffer`, represents a spatially coherent and continuous measure of population density surrounding each node. It provides a realistic proxy for potential mobility demand and is used in subsequent modelling stages to approximate traffic intensity and user presence in the absence of direct traffic counts.

3.6.2 Synthetic Traffic Estimation Using the Radiation Model

To generate realistic origin–destination flows between nodes, we adopted the Radiation Model introduced by Simini et al. [48], which provides a parameter-free and universally applicable framework for modelling human mobility and migration patterns. The model estimates the expected number of trips between two locations based solely on their populations and the distribution of surrounding opportunities, without requiring calibrated friction functions or empirical mobility data.

For each pair of nodes (i, j), we use the previously computed population within the local buffer as a proxy for the population of the origin (m_i) and destination (n_j). The Euclidean distance between nodes defines the radius of influence, within which we compute s_{ij}, the total population contained in all intermediate nodes. These quantities are then combined through the radiation formula:

$$T_{ij} = T \cdot \frac{m_i n_j}{(m_i + s_{ij})(m_i + n_j + s_{ij})},$$

where T is a scaling factor ensuring that the total number of generated trips matches a predefined system-wide trip volume.

This procedure yields a full origin–destination matrix where flows emerge naturally from the spatial distribution of population, producing a realistic and theoretically grounded approximation of mobility demand in the absence of observed traffic data.

3.6.3 Calibration of Temporal Car-Driver Profiles from OD Data

To anchor the synthetic flows generated by the radiation model to observed mobility patterns, we used the official origin–destination (OD) matrix *"Matrice OD2016 – Passeggeri – Scenario 2020"* published by Regione Lombardia [43]. The dataset reports the number of trips between traffic zones, disaggregated by province of origin and destination, mode and role (e.g. driver), and time-of-day bands.

First, we filtered the OD matrix to retain only those records where both the origin and destination provinces belong to Lombardy. We then focused on internal urban mobility in Milan by selecting OD pairs where both the origin and destination zones correspond to the municipality of Milano, after harmonising zone labels (e.g. removing numeric suffixes such as "MILANO 2"). For this subset, we computed (i) the total number of internal trips within Milan and (ii) the total number of trips performed as car drivers, obtained by summing the relevant columns that represent different trip purposes or user categories in the "driver" role.

To derive a temporal profile of car-driver mobility, we grouped the Milan-only car-driver trips by time-of-day band (`FASCIA_ORARIA`) and aggregated the corresponding counts. Dividing these band-specific totals by the overall number of car-driver trips yields a set of percentages that describe the empirical distribution of car-driver movements across the day. The resulting vector of time-of-day shares is used to weight the synthetic flows, ensuring that the generated traffic reflects not only spatial patterns but also realistic temporal variation in urban mobility.

3.6.4 From Synthetic Flows to Electric Vehicle Charging Demand

The radiation model described above provides synthetic origin–destination flows in terms of generic person trips between nodes. To convert these into an estimate of daily electric vehicle (EV) charging demand, we combined three empirical pieces of information:

- the share of trips performed by car drivers within Milan, computed from the OD matrix as the ratio between trips with role "driver" and the total number of trips (approximately 18.4%);
- the share of electric vehicles in the car fleet of Milan, obtained by comparing the number of registered EVs with the total number of cars (around 0.39%);
- the ratio between the observed daily number of charging sessions in the (broader) reference area and the stock of EVs, which provides an estimate of the probability that an EV performs at least one charging session per day (about 8.6% in our case).

Under the simplifying assumption that these proportions are independent and spatially homogeneous within the study area, we scale each synthetic flow T_{ij} by a constant factor that encapsulates: (i) the probability that a trip is made by car, and (ii) the probability that an EV associated with that flow actually charges during the day. The resulting quantity,

$$D_{ij} = \alpha \, T_{ij},$$

defines the expected number of daily EV charging sessions associated with the OD pair (i, j), where α aggregates the empirical percentages (e.g. by approximating the car share as 20% and the charging probability as 10%). In the implementation, this corresponds to adding a column `demand_flow` to the OD table, obtained by rescaling the synthetic traffic values.

To introduce temporal structure, we used the time-of-day distribution of car-driver trips derived from the OD matrix (Section 3.6.3). For each time band h (e.g. hourly interval), we computed a weight w_h representing the share of total car-driver trips occurring in that band, and normalised these weights so that $\sum_h w_h = 1$. We then formed the Cartesian product between the OD pairs and the time bands and distributed the daily demand as

$$D_{ijh} = D_{ij} \, w_h,$$

thus obtaining a time-resolved charging demand at the resolution of the original OD bands.

Finally, to reach a finer temporal granularity suitable for operational analysis, each time band was subdivided into four 15-minute windows. For each (i, j, h), the corresponding quantity D_{ijh} was split into four non-negative values that sum back to D_{ijh}, and assigned to the sub-intervals within that hour. This yields a data set D_{ijhw}, where w indexes the 15-minute windows, providing a realistic spatio-temporal demand profile for EV charging at the network level.

3.6.5 Disaggregation of Demand by Connector Power

The spatio-temporal demand D_{ijhw} described above represents the expected number of daily charging sessions between each OD pair, time window and day, but is still aggregated over all connector types. To obtain a more realistic representation of infrastructure usage, we further disaggregated this demand by connector power classes using empirical information provided by Atlante on existing stations.

From the Atlante dataset, we selected all charging points located in the province of Varese and, for each connector, we considered its rated power and the average daily number of sessions. Connector powers were then grouped into a small number of representative power ranges (e.g. 20–48 kW, 50–75 kW, 90–145 kW, 150–225 kW, 300 kW, and $>$ 300 kW). For each power range, we summed the corresponding average daily sessions and computed the relative share of sessions in that range with respect to the total. Normalising these shares yields a vector of weights

$$\{w_p^{\mathrm{pow}}\}_p, \qquad \sum_p w_p^{\mathrm{pow}} = 1,$$

which can be interpreted as the empirical probability that a generic charging session in the area is served by a connector in power class p.

To assign power information to the synthetic demand, we formed the Cartesian product between the OD–time–window demand table and the set of power classes. For each combination (i, j, h, w, p), we allocated a fraction of the demand proportional to the corresponding weight:

$$D_{ijhw}^{(p)} = D_{ijhw}\, w_p^{\mathrm{pow}}.$$

In the implementation, this corresponds to adding a column `flow_conn_power` obtained by multiplying the window-level demand by the power-class weight. Very small flows were discarded (by applying a minimum threshold before rounding), and the remaining values were rounded to the nearest integer number of sessions. Finally, each record was associated with the start time of the corresponding 15-minute window.

The resulting dataset provides, for each OD pair, time window, and power class, an integer-valued estimate of the number of charging sessions. This representation is rich enough to support subsequent analyses on infrastructure sizing and technology mix, while remaining consistent with both the synthetic mobility model and the empirical usage patterns observed in the Atlante network.

Chapter 4

Preliminary analysis

Before constructing the models, it becomes crucial to understand the interdependencies between charging infrastructure and traffic dynamics. This chapter presents a systematic exploration of EV charging sessions within the context of urban traffic, environmental considerations, and the emergent patterns of electric mobility.

Through a series of correlation analyses, we investigate the relationships between charging sessions and a variety of factors including traffic volume, power output of charging stations, the number of connectors available, and proximity to points of interest. Employing Spearman's Rank [50], Kendall's Tau [25], and Pearson's Correlation Coefficients [41], we unearth subtle nuances in data that reveal the complexities of EV charging behavior.

Temporal patterns of charging sessions are dissected through a week-by-week analysis, uncovering trends that distinguish weekdays from weekends. This temporal scrutiny extends to the evaluation of the uniformity of charging sessions across cities, the consistency of average session durations, and the discernible monthly variations in EV charging activity. The implications of these patterns are discussed in detail, providing a narrative on the efficacy and accessibility of the current charging infrastructure.

Additionally, a regional analysis underscores the correlation between traffic data and EV charging sessions. Here, we delve into the total and mean number of sessions, reflecting on how the penetration rate of electric vehicles within a region influences these figures. This analysis extends to consider the nuances of electric vehicle adoption and how they impact the utilization of charging stations.

4.1 Correlation analysis

As a preliminary step prior to the deployment of advanced machine learning models, we conducted a series of exploratory correlation analyses. The first objective was to identify which features exhibit the strongest statistical associations with the average number of charging sessions per day at individual electric vehicle (EV) charging stations. Specifically, we examined the correlation between daily average charging sessions and factors that are more intuitively related to the usage of a charging station, such as connector power ratings, local traffic intensity, and the number of charging connectors within the station. Then we extended our analysis to various points of interest (POIs). Then, we examined the interdependencies among our key predictors to guide the selection of statistical models best suited for pinpointing the most advantageous locations for new charging stations. To quantify these relationships, we employed three classical correlation measures: Pearson's r, Spearman's rho, and Kendall's tau.

4.1.1 Correlation between Annual traffic, Charging Station Power Output, and Number of Charging Points and Daily Charging Sessions

To select the most appropriate correlation measure for relating annual block traffic to charging station attributes, we first applied Shapiro-Wilk tests [46] to each predictor and to the response variable (mean daily sessions) to assess normality. We then calculated Pearson's r, Spearman's ρ, and Kendall's τ for three primary features—annual block traffic, station power capacity, and number of charging points. Table 4.1 presents the full results, including each normality test outcome, correlation estimates, and our recommended metric. The last column indicates the recommended metric based on distributional assumptions and observed coefficients.

Table 4.1: Normality Tests and Correlation Coefficients with Suggested Metric

Feature	**Normality**	**Pearson's r**	**Spearman's ρ**	**Kendall's τ**	**Suggested**
block_traffic	0.0	0.1885	0.2660	0.1801	Spearman
conn_power	0.0	0.6276	0.2870	0.2109	Spearman
n_cp	0.0	0.4893	0.4188	0.3328	Spearman

All three candidate predictors violated the univariate normality assumption (Shapiro–Wilk $p = 0.0$ in each case; see Table 4.1). Consequently, reliance on Pearson's r alone risks distorted inference due to skewness, kurtosis, and outliers. Below, we examine each feature in turn and justify the choice of Spearman's ρ as the primary association metric.

1. Annual Block Traffic (`block_traffic`)

- *Pearson's* $r = 0.1885$ suggests a weak linear trend, but the non-normal, right-skewed distribution of annual traffic counts—and a handful of extremely busy hexagons—inflate variance and bias r downward.
- *Spearman's* $\rho = 0.2660$ more faithfully captures the underlying monotonic increase in charging sessions as traffic grows, free from the undue influence of outliers.
- *Kendall's* $\tau = 0.1801$ corroborates the general ordering but is less sensitive in our large sample.

The stronger rank-based signal indicates nonlinearity (e.g. diminishing returns at extreme traffic volumes) and justifies using Spearman's ρ to report the block-traffic-charging relationship.

2. Connector Power Output (`conn_power`)

- *Pearson's* $r = 0.6276$ implies a substantial linear association, yet connector power takes on only a few discrete levels (e.g. 22 kW, 50 kW, 150 kW, etc...), yielding many tied values and potential "step" effects.
- *Spearman's* $\rho = 0.2870$ and *Kendall's* $\tau = 0.2109$ are markedly lower, reflecting that relative station rankings by power do not translate uniformly into session counts—usage spikes at ultra-fast levels rather than scaling smoothly.

Although Pearson's r is numerically large, its assumptions are violated by discrete, heteroskedastic data. Spearman's ρ better balances sensitivity to monotonic trends with robustness to ties and outliers, making it the preferred metric.

3. Number of Charging Points (`n_cp`)

- *Pearson's* $r = 0.4893$ indicates a moderate linear effect, but station sizes cluster (e.g. 2, 4, 6 points), again creating ties and skew.
- *Spearman's* $\rho = 0.4188$ remains substantial, demonstrating a clear rank-order increase in sessions with more points, albeit with diminishing marginal gains.
- *Kendall's* $\tau = 0.3328$ is consistent with Spearman but slightly lower in absolute magnitude.

The monotonic relationship is strong and consistent; Spearman's ρ offers a more reliable, assumption-light summary than Pearson.

4.1.2 Comparison of Old and New Feature Groupings

This section presents a direct comparison between the "old" and "new" POI groupings by reporting their peak correlation coefficients (Pearson's r, Spearman's ρ, and Kendall's τ) with our target variable, average daily charging sessions. The purpose of this analysis is to illustrate how aggregating granular features into broader categories (1) strengthens the observed signal, (2) mitigates noise and the impact of tied values, and (3) produces a more parsimonious, interpretable feature set for subsequent demand modeling. Tables 4.2 and 4.3 list the full spectrum of correlation coefficients for the deviation-flow–based features (old grouping) and the node-serving–based categories (new grouping), respectively.

As Table 4.3 (and its counterpart for the old grouping) shows, no single POI category exhibits a strong bivariate correlation ($|r|>0.5$) with average daily charging sessions. Even

Table 4.2: Correlation coefficients between old grouping features and average daily charging sessions

Feature	**Pearson's r**	**Spearman's ρ**	**Kendall's τ**
mall_foot_walking_5_n	0.0942	0.1655	0.1289
transportation_driving_car_10_n	0.0894	0.1816	0.1253
hospitality_driving_car_10_n	0.0833	0.1673	0.1178
mall_foot_walking_10_n	0.0791	0.1624	0.1228
mall_driving_car_5_n	0.0772	0.1736	0.1245
institutional_driving_car_10_n	0.0748	0.1650	0.1160
transportation_driving_car_5_n	0.0743	0.1766	0.1218
hospitality_foot_walking_10_n	0.0727	0.1039	0.0728
mall_driving_car_10_n	0.0696	0.1691	0.1187
hospitality_driving_car_5_n	0.0664	0.1272	0.0905
transportation_foot_walking_5_n	0.0634	0.1538	0.1061
transportation_foot_walking_10_n	0.0602	0.1625	0.1103
shopping_driving_car_10_n	0.0534	0.1602	0.1125
leisure_driving_car_10_n	0.0496	0.1512	0.1062
institutional_driving_car_5_n	0.0496	0.1169	0.0837
hospitality_foot_walking_5_n	0.0463	0.1068	0.0755
shopping_foot_walking_10_n	0.0452	0.0963	0.0704
shopping_driving_car_5_n	0.0385	0.1265	0.0902
institutional_foot_walking_10_n	0.0248	0.0583	0.0433
shopping_foot_walking_5_n	0.0240	0.0770	0.0586
leisure_driving_car_5_n	0.0222	0.1089	0.0788
leisure_foot_walking_10_n	0.0106	0.0645	0.0470
logistic_driving_car_10_n	0.0049	0.0617	0.0495
logistic_driving_car_5_n	-0.0063	0.0247	0.0200
leisure_foot_walking_5_n	-0.0152	0.0387	0.0284
institutional_foot_walking_5_n	-0.0173	0.0144	0.0113
logistic_foot_walking_10_n	-0.0175	-0.0168	-0.0137
logistic_foot_walking_5_n	-0.0239	-0.0250	-0.0204

Table 4.3: Correlation coefficients between new grouping features and average daily charging sessions

Feature	Pearson's r	Spearman's ρ	Kendall's τ
fuel	0.1912	0.1873	0.1440
mall	0.1167	0.1595	0.1292
fast_food	0.0924	0.1788	0.1402
sustenance_long	0.0817	0.1383	0.1018
department_store	0.0792	0.0677	0.0552
furniture	0.0682	0.1079	0.0868
parking	0.0547	0.1056	0.0725
transportation	0.0543	0.1050	0.0805
amenity_facilities	0.0500	0.1172	0.0870
wholesale	0.0458	0.0381	0.0311
office	0.0344	0.0612	0.0472
shop	0.0274	0.0967	0.0697
amenity_others	0.0244	0.0637	0.0505
sustenance_short	0.0182	0.0592	0.0433
financial	0.0162	0.0405	0.0315
entertainment	0.0131	0.0664	0.0518
leisure	0.0126	0.0550	0.0415
healthcare	0.0079	0.0511	0.0395
house	0.0069	0.0502	0.0363
sports	-0.0103	-0.0042	-0.0024
supermarket	-0.0139	0.0548	0.0429
tourism	-0.0187	0.0394	0.0303
education	-0.0417	-0.0231	-0.0170
spiritual	-0.0569	-0.0302	-0.0233
public_service	-0.0629	-0.0306	-0.0235

the highest Pearson coefficient—0.19 for "fuel"—remains in the weak-to-moderate range, and the top rank-based measures (Spearman's $\rho \approx 0.19$, Kendall's $\tau \approx 0.14$) likewise fall short of conventional thresholds for strong association.

This pattern is not unexpected in complex urban systems, where charging demand arises from multifaceted interactions among land-use, mobility, and socioeconomic factors. The absence of dominant univariate predictors underscores the necessity of multivariate modeling frameworks capable of integrating many weak signals. In particular, tree-based ensemble methods (Random Forest, XGBoost) are well suited to this setting: they harness the combined explanatory power of correlated, low-signal features through hierarchical splitting and aggregation, capture non-linear interactions, and remain robust to multicollinearity. By leveraging the full feature set rather than relying on any individual POI category, these models can uncover the subtle, joint effects that drive charging behavior.

4.1.3 Model Selection Informed by POI Correlation Structure

To determine the most appropriate predictive model for forecasting charging session counts, we first analyzed the linear interdependencies among our aggregated POI feature set using Pearson's correlation coefficient, which is well-suited for buffer-aggregated counts and provides a clear measure of pairwise linear association. In the deviation-flow approach, the various buffer-mode combinations inherently derive from the same underlying count data at different spatial scales, so high correlations among these features are expected and need not be re-examined formally. By contrast, our node-serving categories represent distinct POI aggregations whose interrelationships are not predetermined by construction. Therefore, we explicitly computed and analyzed the Pearson correlation matrix for the node-serving features (Figure 4.6) to identify non-trivial associations—such as between “shop” and “sustenance” groups—that could impact model performance and guide our choice of robust, multicollinearity-resistant algorithms.

These observed inter-category correlations—indicative of multicollinearity—confirm the goodness of the use of ensemble tree-based methods. Random Forests address collinearity through bootstrap aggregation and randomized feature selection at each split, thereby reducing variance and preventing any single correlated group from dominating the model. XGBoost further builds on this robustness by fitting residuals iteratively with regularized decision trees, automatically down-weighting redundant predictors and capturing complex, non-linear interactions without requiring manual de-correlation or dimensionality reduction. Consequently, we adopt Random Forest [10] and XGBoost [12] as our primary modeling approaches, leveraging their inherent resistance to multicollinearity and strong performance in the presence of correlated, high-dimensional feature spaces.

4.2 Temporal Analysis of Charging Sessions

Understanding the hourly patterns of electric-vehicle charging is essential for designing resilient and user-focused charging networks. By comparing weekday and weekend session counts and durations, we can reveal distinct behavioral rhythms—such as morning commuter peaks versus leisure-driven usage—that inform station siting, load-balancing strategies, and staffing levels. Moreover, these insights provide a robust empirical foundation for generating realistic synthetic demand traces, enabling stress-testing of control algorithms and expansion scenarios underpinned by high-quality, real-world data.

The study utilizes comprehensive data on charging sessions across various regions, analyzing the frequency of sessions by hour over the course of a typical week. This temporal framing enables a detailed examination of patterns, highlighting both the consistency and variability in charging behavior.

The ensuing sections present a comprehensive analysis of the charging session data, starting with an examination of weekday patterns, followed by a comparative analysis of weekend behavior. The discussion encapsulates the implications of these patterns for charging infrastructure development and offers insights into potential strategies for accommodating the observed demand fluctuations.

Table 4.4: Average number and duration of charging sessions by hour on a random weekday and weekend

Hour	**# Sessions**		**Avg. Duration (min)**	
	Weekday	**Weekend**	**Weekday**	**Weekend**
0	12 790	7 934	66.49	57.71
1	9 561	5 794	62.96	55.39
2	9 007	5 354	67.92	58.13
3	12 042	5 974	57.34	54.12
4	24 830	8 826	51.12	51.32
5	54 589	17 744	49.94	50.40
6	104 506	33 647	52.72	49.38
7	152 726	57 741	55.96	51.52
8	169 286	82 233	57.83	51.56
9	181 017	97 614	56.66	52.74
10	192 360	97 106	55.00	53.01
11	193 406	82 096	56.37	52.83
12	179 432	72 364	59.92	54.53
13	176 890	77 083	52.81	50.27
14	188 902	84 527	52.15	50.37
15	200 063	90 262	51.93	50.01
16	207 765	88 987	51.16	48.76
17	188 493	75 424	49.84	50.84
18	138 834	54 355	51.13	48.80
19	92 825	37 826	52.81	50.85
20	65 627	28 426	52.36	49.69
21	48 198	21 885	55.84	50.77
22	32 996	15 459	54.93	50.73
23	21 483	10 005	63.83	54.06

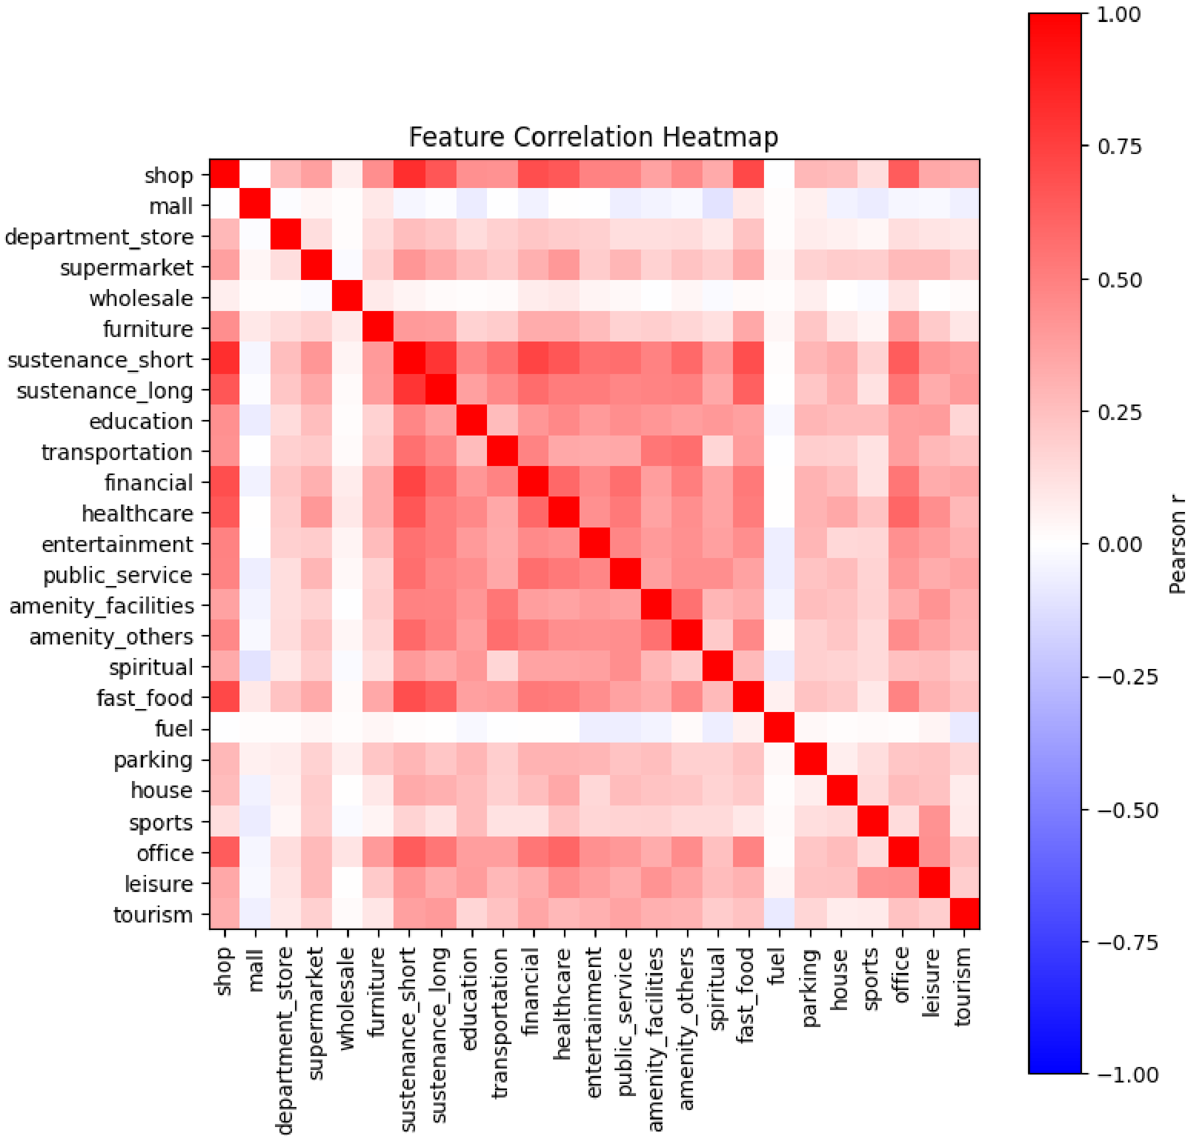

Figure 4.1: Correlation heat map among Point of Interest Groups

4.2.1 Weekdays Analysis of Charging Sessions

The data indicate a distinct daily pattern in the utilization of EV charging stations. Early morning hours (from midnight to 5 AM) exhibit a gradual increase in charging sessions, with the lowest counts occurring at 23:00 (21,483 sessions) and steadily rising to a peak at 05:00 (54,589 sessions). This suggests that a significant number of EV users may prefer to charge their vehicles overnight, potentially to take advantage of lower electricity rates or to ensure a full charge by morning.

The usage escalates sharply in the morning, with the highest number of sessions observed during typical commuting hours. A pronounced peak is reached between 08:00 and 09:00, aligning with conventional morning rush hours. The number of sessions remains high throughout the day, peaking at 11:00 with 193,406 sessions, indicating sustained demand for charging during the working hours.

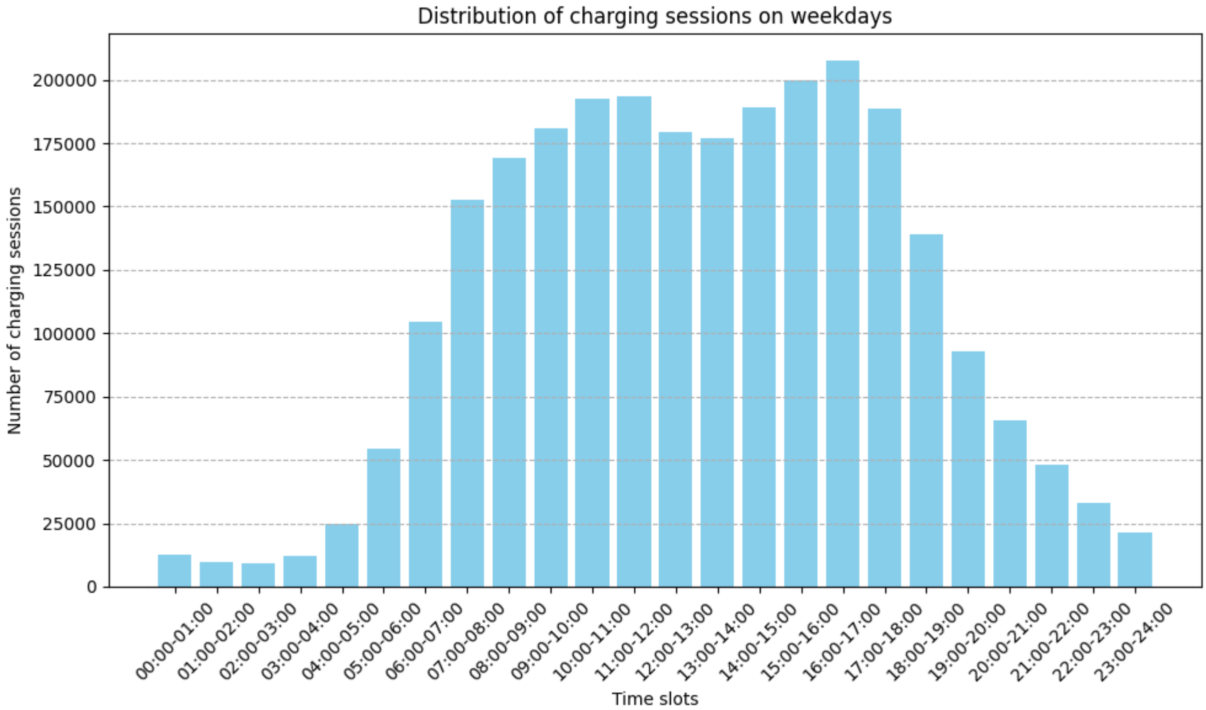

Figure 4.2: Weekdays Analysis of Charging Sessions

Interestingly, there is a gradual decline in sessions post-midday, with a slight plateau between 14:00 and 16:00, followed by a more substantial decrease after 17:00. The period from 18:00 to midnight shows a significant drop-off in charging sessions, which could reflect the end of the workday and less demand for public charging as individuals return home.

The observed distribution suggests that charging station usage is heavily influenced by typical workday patterns, with stations experiencing higher demand during commuting times and standard working hours.

4.2.2 Weekend Analysis of Charging Sessions

Building upon our understanding of EV charging sessions' temporal distribution, this section focuses on analyzing charging patterns during weekends. Such an analysis is vital for identifying divergences from weekday charging behavior, which may influence operational and strategic decisions differently.

The weekend data reveals a distinct pattern compared to weekdays, with a notable shift in charging session peaks and a general reduction in the total number of sessions. The earliest hours (00:00 to 05:00) see a gradual increase in charging activities, starting from 7,934 sessions at midnight and peaking at 17,744 by 05:00. This pattern is consistent with the overnight charging trend observed during weekdays, albeit with lower overall numbers, suggesting a continued preference for nighttime charging among EV users, possibly due to similar motivations such as lower rates and the assurance of a fully charged vehicle by morning.

However, the morning rise in charging sessions starts later and ascends more sharply,

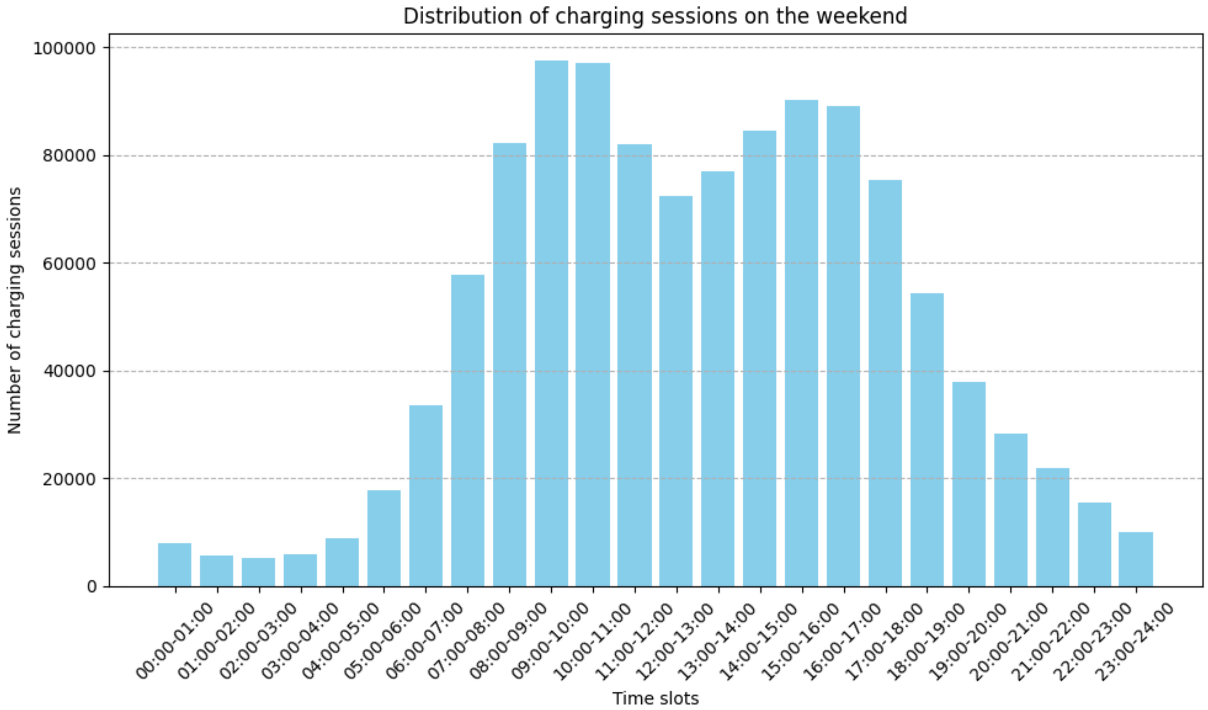

Figure 4.3: Weekend Analysis of Charging Sessions

with a significant increase leading to a peak at 09:00 (97,614 sessions), which is slightly later than the peak times observed during weekdays. The numbers remain high around midday but do not exhibit the same level of sustained peak as seen on weekdays. Notably, the highest number of sessions occurs at 10:00 (97,106 sessions), with a gradual decrease thereafter.

Afternoon hours continue to show robust charging activity, with a gradual decline setting in after 15:00. The evening sees a faster reduction in sessions compared to weekdays, with numbers dropping significantly after 18:00, leading to a quieter night.

The weekend charging pattern highlights a shift towards later morning peaks and a more pronounced midday decline in charging sessions. This shift indicates a potential adjustment in user behavior, possibly reflecting less structured weekend schedules or leisure-related travel that does not conform to the regular workday patterns.

4.2.3 Uniformity Across Cities in Total Charging Sessions

An intriguing aspect of our analysis is the consistency in the number of total charging sessions across various cities. To formally assess this uniformity, we conducted chi-square goodness-of-fit tests comparing each city's observed hourly counts against the overall expected distribution. In every case, the resulting p-values exceeded conventional significance thresholds (e.g. $p > 0.05$), indicating that no city's temporal profile deviates significantly from the common pattern. This indicates that despite the geographical and socio-economic differences that might influence EV usage and charging needs, the data reveals a uniform pattern in the distribution of charging sessions throughout the day for both weekdays and weekends. This uniformity suggests a widespread, homogeneous behavior among EV users in different urban settings. The lack of substantial variance was further validated through chi-square tests, which confirmed that the differences in the number of charging sessions across cities are not statistically significant. This finding implies that the observed temporal patterns of charging behavior are likely reflective of broader, perhaps national-level, trends in EV charging habits rather than being specific to local idiosyncrasies.

4.3 Analysis of Average Charging Session Duration by Hour

This analysis aims to elucidate the temporal patterns of EV charging session durations across different times of the day, with a special focus on distinguishing between weekdays and weekends. Such differentiation is crucial, as it accommodates the variance in user behavior and infrastructure demand, laying the groundwork for targeted improvements in charging station management and planning.

As done in the previous section, this analysis is bifurcated into two distinct temporal segments—weekdays and weekends—to capture the variability in charging behaviors that may arise due to the typical weekly schedules of users. This segmentation allows for a nuanced understanding of charging patterns, facilitating a more tailored approach to infrastructure development and policy making.

4.3.1 Weekdays analysis

This section examines how the duration of these sessions fluctuates throughout the 24-hour cycle on weekdays, providing insights into user charging habits and the operational dynamics of charging stations.

The data presents a nuanced view of charging session durations across different hours of the day. Notably, the average duration of charging sessions tends to be longer during late-night to early-morning hours, peaking at 02:00 with an average of 67.923 minutes. This could be indicative of overnight charging, where users are less pressed by time constraints and possibly take advantage of off-peak electricity rates.

A gradual decrease in session duration is observed from the early morning hours, reaching its shortest average duration of 49.841 minutes at 05:00. This trend reverses as the morning progresses, with durations extending until 08:00, likely reflecting a mix of early commuters starting their day and overnight sessions concluding.

The daytime hours from 09:00 to 17:00 show relatively shorter and more stable session durations, with averages hovering around 51 to 57 minutes. This period likely represents a combination of opportunity charging by commuters and daytime users managing their daily activities.

Interestingly, a second increase in session duration occurs in the late evening, starting

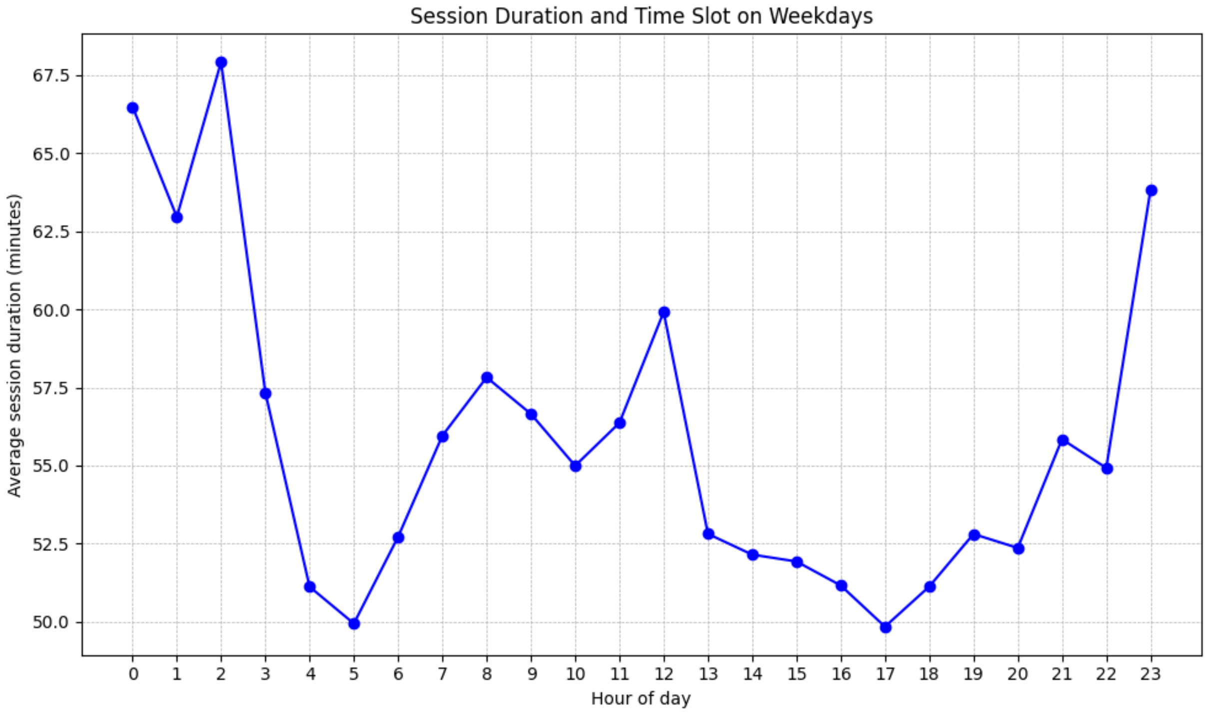

Figure 4.4: Average Charging Session Duration by Hour for Weekdays

at 21:00 and peaking again at 23:00 with an average duration of 63.833 minutes. This evening peak suggests a pattern of users plugging in their vehicles for overnight charging before the end of the day, potentially to ensure a full charge for the next day.

The extended durations during the night and early morning hours emphasize the need for adequate overnight charging facilities, particularly in residential areas. Meanwhile, the relatively consistent and shorter durations during the day highlight the importance of providing accessible, fast-charging options to accommodate the needs of daytime users.

4.3.2 Weekend Analysis of Average Charging Session Duration

Following the weekday analysis, this section delves into the average duration of electric vehicle (EV) charging sessions during weekends.

The weekend data reveals a distinctive pattern compared to weekdays, with some similarities in the late night and early morning hours but notable differences during the day. The longest sessions occur just after midnight, at 00:00, with an average duration of 57.711 minutes, decreasing slightly through the early morning hours. This might reflect a continuation of the preference for overnight charging, albeit with slightly shorter durations than those observed during weekdays.

A steady decrease in session duration is seen from the early hours of the morning, reaching its lowest at 16:00 with an average of 48.755 minutes. This contrasts with the weekday pattern, where session durations began to increase in the early evening. The shorter durations throughout the day suggest that weekend charging sessions may be more oriented towards quick, opportunistic charges rather than prolonged sessions.

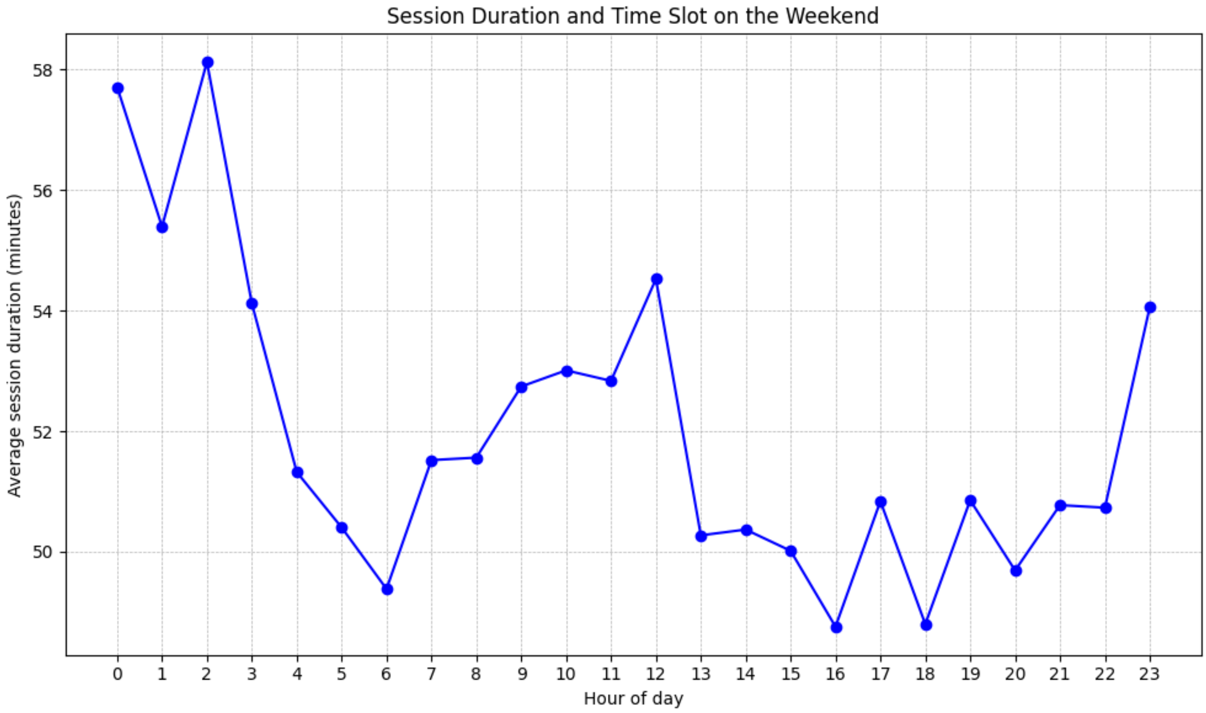

Figure 4.5: Average Charging Session Duration by Hour for Weekends

Notably, the average duration begins to increase again in the late evening, reaching a second peak at 23:00 with an average duration of 54.059 minutes. This late evening increase might be attributed to users preparing for the week ahead or taking advantage of off-peak rates before the new week begins.

Not different from what is concluded for the weekdays, the data suggests that there is a consistent demand for overnight charging facilities, but also a need for efficient, quick-charge options during daytime hours to accommodate the shorter, more opportunistic charging sessions prevalent on weekends.

4.3.3 Consistency in Average Session Durations Across Cities

Similar to what we found during the number of sessions, the analysis of average session durations reveals a consistent pattern across different cities. This consistency was observed in both the peak hours for charging sessions and the average duration of these sessions, with no significant variations detected among the cities included in the study. Chi-square tests were utilized to assess the variability in average session durations across cities, yielding results that support the hypothesis of uniformity. These outcomes indicate that the average time EV users spend charging their vehicles does not significantly differ by location, reinforcing the notion of standardized charging behavior. The implications of this uniformity are profound, suggesting that strategies for infrastructure improvement and policy development can be applied broadly, with expectations of similar user responses across diverse urban landscapes.

4.4 Monthly Trends in EV Charging Sessions

This section examines the monthly distribution of electric vehicle (EV) charging sessions over a two-year period. Understanding these trends is essential for infrastructure planning and management, as it enables the anticipation of demand fluctuations and the strategic allocation of resources.

4.4.1 Analysis of Monthly Charging Session Trends

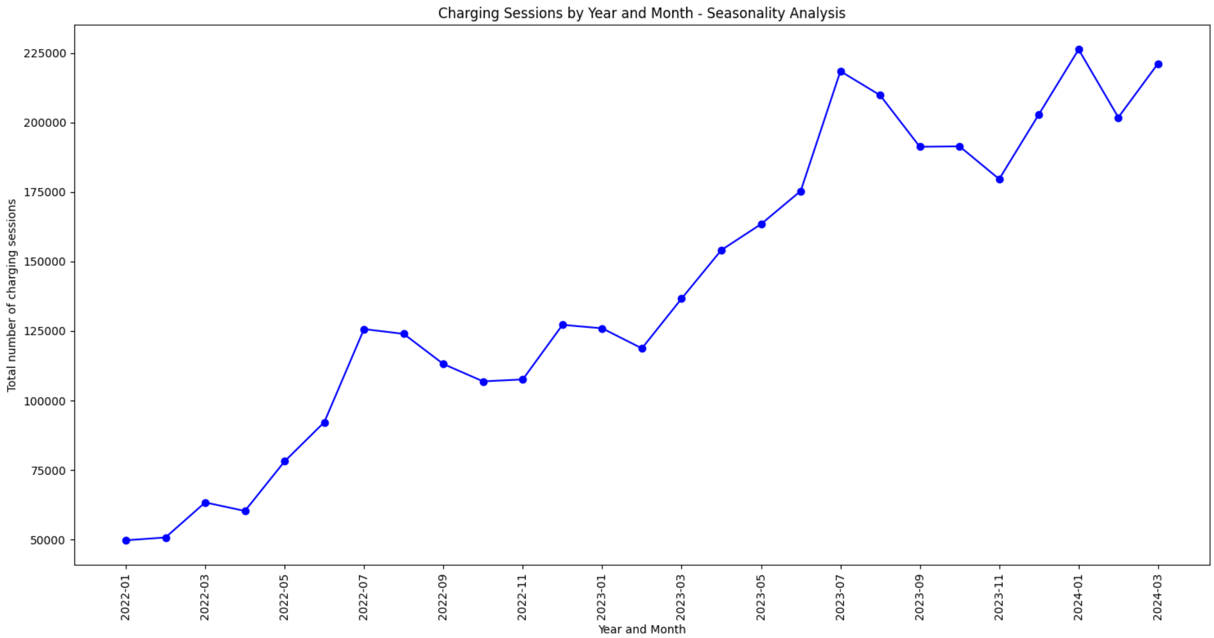

Figure 4.6: Charging Sessions by Year and Month - Seasonality Analysis

The data presents a clear ascending trend in the total number of EV charging sessions, demonstrating a growing adoption and utilization of EVs. From an initial count of 49,767 sessions in January 2022, there is a consistent increase in monthly session numbers, reaching a peak of 226,251 sessions by January 2024. This upward trajectory underscores the expanding footprint of EVs and the corresponding need for charging infrastructure.

Notably, the analysis reveals two significant peaks in charging sessions, occurring in July and December/January of each year. The July peak aligns with the summer holiday period, suggesting a surge in travel-related charging as individuals and families embark on vacations and road trips. Similarly, the December/January peak corresponds with the winter holiday season, a time traditionally associated with increased travel and, consequently, a heightened demand for EV charging services. These seasonal spikes reflect the influence of holiday travel patterns on EV charging behavior, emphasizing the need for infrastructure readiness to accommodate these periodic surges in demand.

The observed monthly trends and seasonal peaks have profound implications for stakeholders in the EV ecosystem. For infrastructure operators and utility providers, these patterns highlight the importance of dynamic capacity planning to ensure that the charging network can accommodate varying levels of demand throughout the year. Additionally, the data can inform targeted marketing and pricing strategies, encouraging charging during off-peak periods to balance the network load.

4.4.2 Regional Analysis of EV Charging Session Correlations

The transition towards electric vehicles (EVs) as a cornerstone of sustainable transportation requires not only a national vision but also a nuanced understanding of regional dynamics. Italy, with its diverse geographical landscapes, economic activities, and urbanization levels, presents a unique case study for examining the proliferation of EV charging infrastructure and its utilization. This analysis aims to uncover the regional intricacies that influence EV charging session trends across the country. By investigating the Pearson correlation coefficients for monthly EV charging sessions against a backdrop of national trends, this study offers insights into how regional characteristics sway the adoption and use of EV charging facilities. From the industrialized and urbanized north to the rural landscapes and coastal areas of the south, each region contributes distinct patterns to the national narrative of EV integration. This introduction lays the groundwork for a detailed exploration of these patterns, highlighting the importance of regional factors in shaping Italy's journey towards electrified mobility. Understanding these regional differences is not just academic; it provides actionable insights for policymakers, infrastructure developers, and stakeholders in the EV ecosystem to tailor strategies that address specific regional needs and challenges.

The comprehensive regional analysis of Pearson correlation coefficients between the monthly trends of EV charging sessions and the overall trend across Italy presents intriguing insights:

- Lombardia, with the largest dataset of 120 cities, exhibits a mean correlation coefficient of -0.14, with a substantial deviation (max at 0.20 and min at -0.85). This suggests a broad spectrum of charging behavior, likely influenced by the diverse economic and geographic landscape of the region.

- Piemonte, featuring 40 cities in the analysis, shows a nearly neutral mean correlation of -0.01, but with correlations stretching from -0.43 to 0.20. The variance implies a significant disparity in charging session growth rates, which may reflect varying levels of EV adoption or charging infrastructure availability.

- In contrast, regions with a smaller number of cities analyzed, such as Calabria and Basilicata, each represented by just one city, display positive correlations (0.17 and 0.14, respectively). This positive correlation could indicate more uniform adoption of EVs or consistent growth in charging sessions, potentially driven by focused initiatives or homogeneous user behavior within these regions.

Table 4.5: Regional Statistics of Pearson Correlation Coefficients for EV Charging Sessions.

Region	# of Cities	Mean	Max	Min	Std. Dev.
Lombardia	120	−0.14	0.20	−0.85	0.24
Piemonte	40	−0.01	0.20	−0.43	0.17
Veneto	20	−0.01	0.18	−0.55	0.20
Emilia-Romagna	21	−0.12	0.18	−0.59	0.20
Lazio	15	−0.15	0.14	−0.71	0.25
Campania	14	−0.04	0.20	−0.60	0.22
Trentino-Alto Adige	13	−0.03	0.18	−0.53	0.20
Marche	11	−0.14	0.16	−0.60	0.22
Toscana	6	−0.14	0.18	−0.63	0.30
Liguria	5	0.12	0.17	−0.01	0.07
Sicilia	4	0.04	0.16	−0.24	0.16
Friuli Venezia Giulia	3	−0.12	0.13	−0.38	0.21
Abruzzo	3	0.01	0.08	−0.07	0.06
Puglia	3	−0.14	0.15	−0.50	0.27
Umbria	3	0.04	0.19	−0.27	0.21
Calabria	1	0.17	0.17	0.17	0.00
Basilicata	1	0.14	0.14	0.14	0.00
Sardegna	1	0.10	0.10	0.10	0.00
Valle d'Aosta	1	0.10	0.10	0.10	0.00

- The analysis also highlights regions like Toscana and Campania, with mean correlations of -0.14 and -0.04, respectively. The negative mean suggests a divergence from the national trend in these areas, possibly due to regional specificities in EV usage or infrastructural development lagging behind or advancing differently from the national pace.

- Interestingly, the Liguria region stands out with a positive mean correlation of 0.12 across five cities, indicating a closer alignment with the national EV charging session trend. This could suggest effective regional policies or a more engaged EV community.

These findings underline the complex interplay between regional characteristics and EV charging session trends. The wide range of correlations across Italy suggests that local factors significantly influence EV charging behaviors, requiring region-specific strategies for infrastructure development and policy planning.

This detailed regional analysis sheds light on the intricate dynamics between regional peculiarities and the evolution of electric vehicle (EV) charging sessions across Italy. A closer examination of the Pearson correlation coefficients offers further insights:

- The broad range of correlation values in regions with a high number of cities, such as Lombardia and Piemonte, underscores the complexity of EV adoption and infrastructure utilization within these regions. It suggests that within a single region, cities can exhibit markedly different trends in EV charging sessions, possibly influenced by urban versus rural settings, availability of charging infrastructure, and local economic conditions.

- The positive correlations observed in regions with fewer cities analyzed, such as Calabria and Basilicata, might point towards a more cohesive regional approach to EV infrastructure development or a unified response from the EV user community. These regions might be benefiting from targeted policies or investments aimed at boosting EV adoption and infrastructure, leading to a trend that closely mirrors or even exceeds the national average.
- Regions like Toscana and Campania, displaying negative mean correlations, could be experiencing unique challenges that decelerate the alignment with the national growth trend in EV charging sessions. These challenges could range from infrastructural deficits, such as a lack of adequate charging stations, to cultural and economic factors that might influence the rate of EV adoption.
- Liguria's alignment with the national trend through a positive mean correlation highlights the potential impact of coastal geography on EV charging behavior. Coastal regions might exhibit distinct seasonal trends, with peaks in tourist activity possibly driving a more pronounced increase in EV charging sessions during certain periods.
- The variance in correlations, especially in regions with a significant number of cities included in the study, calls attention to the potential role of regional policy differences. It suggests that localized strategies for promoting EV adoption and expanding charging infrastructure could have a pronounced impact on regional trends in EV charging sessions.
- The analysis also raises questions about the influence of economic factors on EV adoption and infrastructure use. Regions with higher economic activity and greater urbanization may see more rapid increases in charging sessions, which could explain some of the variances observed between regions.

The insights gleaned from this analysis highlight the multifaceted nature of EV charging session trends across Italy. They underscore the need for nuanced, region-specific approaches in policy and infrastructure development to support the continued growth of EV adoption and use. Understanding these regional differences is crucial for tailoring interventions and investments to maximize their impact on the national push towards sustainable transportation.

4.5 Traffic Data Analysis

Understanding the dynamics of EV charging sessions in relation to traffic volume is pivotal for optimizing the placement and utilization of charging stations. This analysis aims to unravel the intricate relationship between traffic patterns and EV charging behaviors, providing insights that could potentially steer policy-making and infrastructural development.

Initially, our investigation focused on the aggregate number of charging sessions occurring within various traffic blocks, seeking patterns that might elucidate the preferences or necessities driving EV users to charge at specific locations. Subsequent to this, we delved into the average charging sessions per traffic block, which shed light on the consistency of charging station usage across different traffic volumes.

To add a layer of depth to our analysis, we factored in the ratio of electric vehicles, leveraging it as a normalization parameter. This ratio serves as a proxy for the penetration of electric vehicles in the region, allowing us to calibrate our findings based on the prevalence of EVs. Through this normalization, we intended to discern whether the frequency of charging sessions was truly reflective of EV usage intensity or skewed by other variables.

The findings of this multifaceted analysis not only contribute to our understanding of current EV charging trends but also help anticipate the evolution of electric mobility within urban centers. The interplay of traffic density and EV charging points narrates a story of current infrastructure and its adaptability to support an electrified future of transportation.

4.5.1 Total number of sessions

The comprehensive analysis of traffic data undertaken in this study was structured to assess the relationship between traffic volume and session occurrences at charging stations. Block traffic, defined as the annual vehicular flow through a specific location, was taken as a primary indicator of traffic intensity.

To initiate this investigation, the initial step involved aggregating all session instances by their corresponding latitude and longitude. This spatial consolidation allowed for a nuanced understanding of session distributions geographically. Subsequently, sessions with identical block traffic values were summed to discern patterns related to traffic frequency.

The resultant graph, plotted from this collated data in 4.7, suggests a non-linear and non-correlated relationship between traffic density and the number of sessions executed at the corresponding charging column. Notably, several blocks with high traffic volumes exhibited a disproportionately low number of recharging sessions. This phenomenon could potentially be attributed to the scant penetration of electric vehicles within these regions. Conversely, less-trafficked blocks showed a surprisingly high session count, possibly indicating the presence of multiple charging stations that necessitate deliberate routing deviations for access.

Impact of Electric Vehicle Penetration on Sum of Charging Sessions Incorporating the ratio of electric vehicles into our analysis provided an additional layer of understanding to the charging session data. By normalizing the session data against the

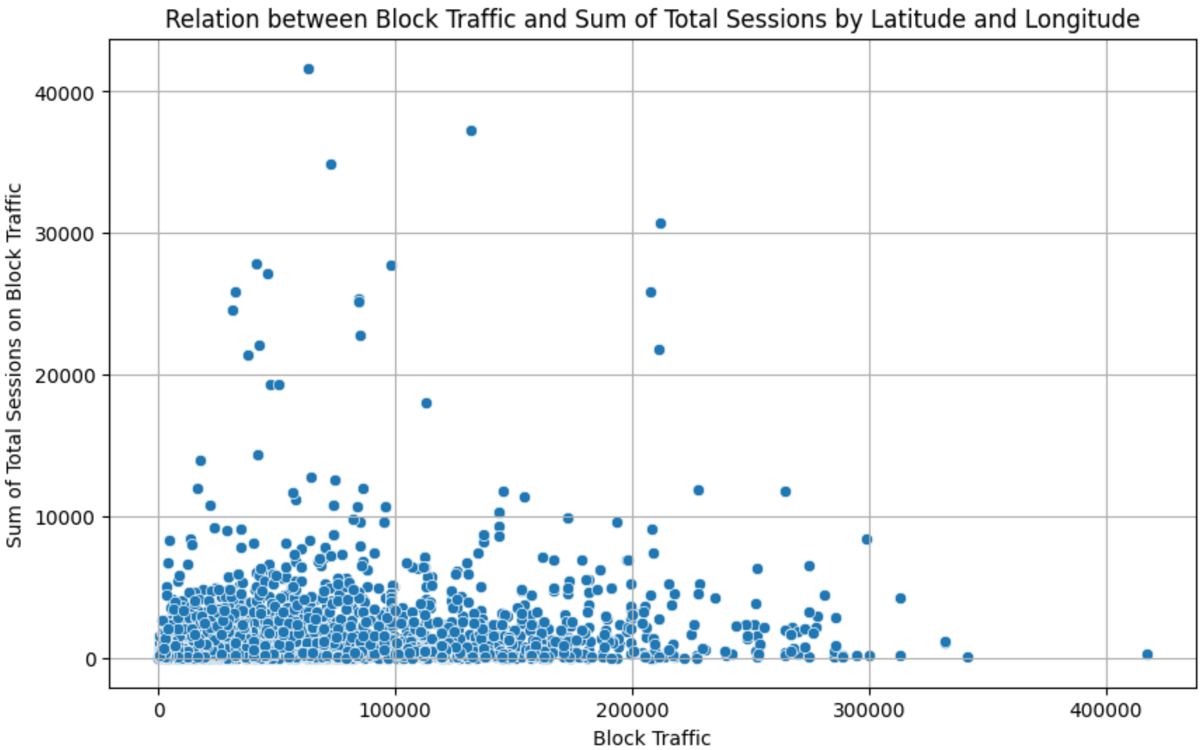

Figure 4.7: Relation between Block Traffic and Sum of Total Sessions

electric car penetration within each block, we aimed to account for the effect of electric vehicle adoption on charging behavior. Surprisingly, as shown in 4.8, the resultant distribution, after applying this electric vehicle-centric normalization, indicated a slight increase in outliers. This suggests that even after adjusting for vehicle type prevalence, there are anomalies in the data where certain traffic blocks have disproportionately high or low numbers of charging sessions relative to the number of electric vehicles present.

4.5.2 Mean number of sessions

A further dimension to our analysis involved computing the average number of recharging sessions per block of traffic. This approach aimed to refine the data by normalizing session counts with respect to traffic volume, thereby reducing the impact of outliers and enabling a clearer interpretation of trends.

Upon calculating these averages, the results corroborated the initial findings to a considerable extent. As shown in 4.9, the modification resulted in a noticeable attenuation of outliers, particularly in blocks registering low traffic yet high session counts. This adjustment can be interpreted as a mitigation of data skewness caused by areas with a high concentration of charging stations or preferential charging locations that may not be directly linked to traffic volume.

However, it is pivotal to note that blocks characterized by high traffic with few recharging sessions exhibited minimal changes in their behavior post normalization. This consistency suggests that while the average may smooth some disparities, the fundamental

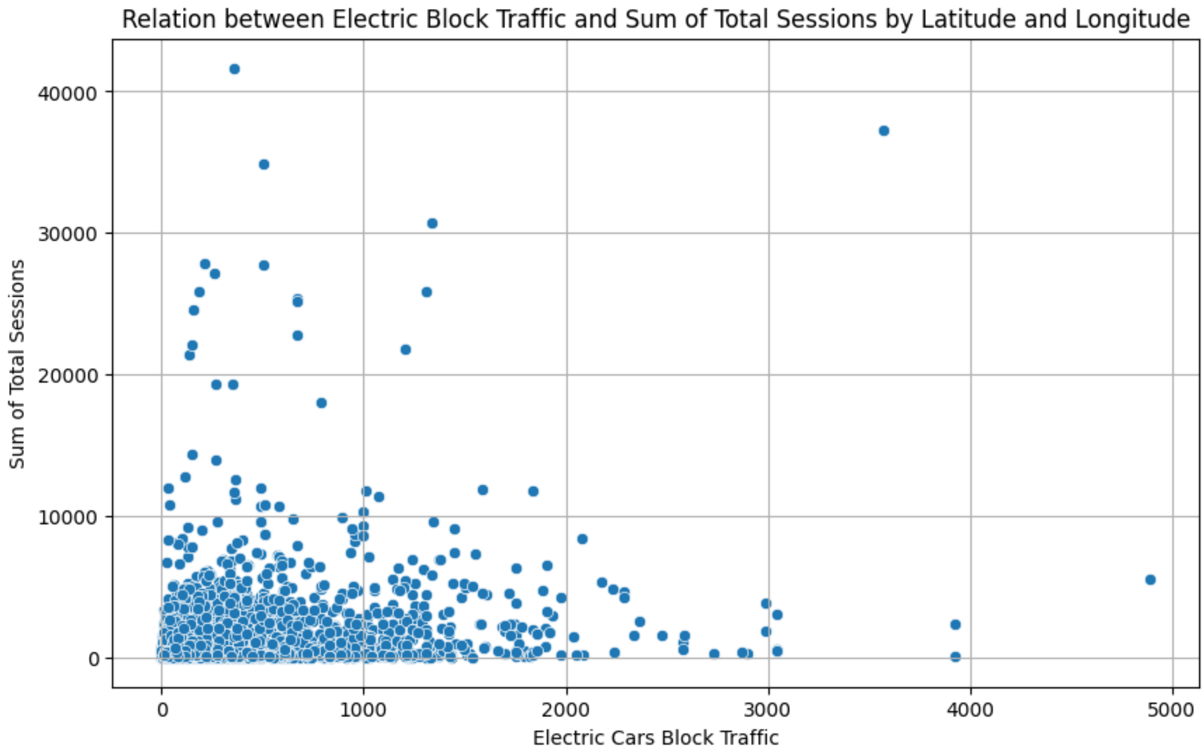

Figure 4.8: Relation between Electric Block Traffic and Sum of Total Sessions

patterns indicating low electric vehicle uptake or charging station utilization in high-traffic areas remain predominantly unchanged.

Influence of Electric Vehicle Penetration on Average Charging Sessions The same normalization process was then applied to the average charging sessions per block traffic, with the intent of discerning more subtle trends in the data. The results, plotted in 4.10, were somewhat unexpected—while the general pattern remained similar to that of the unadjusted averages, the normalization process accentuated the presence of outliers. These deviations were more pronounced, indicating that there are specific blocks where charging sessions per electric vehicle do not align with the overall trend.

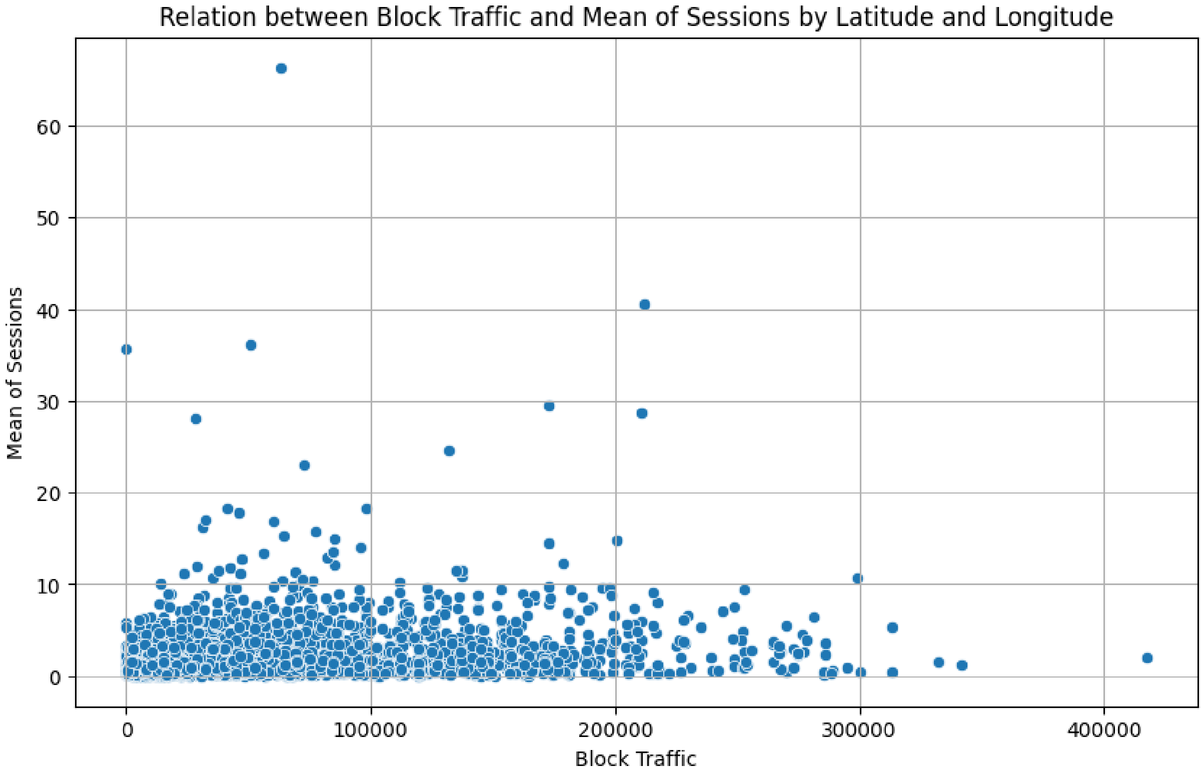

Figure 4.9: Relation between Block Traffic and Mean of Sessions

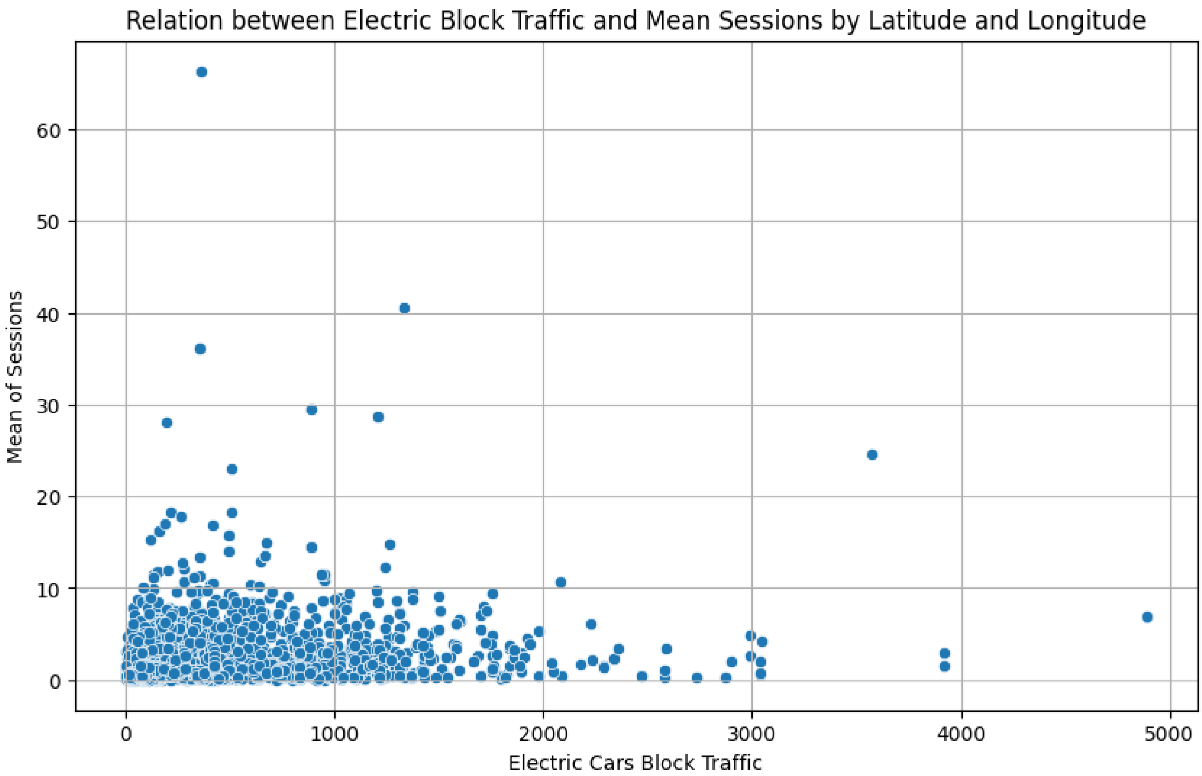

Figure 4.10: Relation between Electric Block Traffic and Mean Sessions

Chapter 5

Data-Driven Insights into Electric Vehicle Charging Stations

This study employs a comprehensive data-driven approach to identify areas where new electric vehicle (EV) charging stations could be most profitable and strategically impactful. The objective is not only to forecast where charging stations are likely to exhibit high utilization, but also to understand the underlying spatial and socio-economic dynamics that drive charging behavior. These dynamics include local demographic characteristics, proximity to major roads and commercial hubs, and—most critically—the presence and composition of nearby points of interest (POIs). Since EV charging requires a non-negligible dwell time, users often engage in other activities while their vehicle charges; therefore, the surrounding urban environment becomes a fundamental determinant of station attractiveness. Insights derived from this analysis support stakeholders such as city planners, policymakers, and private charging infrastructure operators in making informed, evidence-based decisions about network expansion.

To achieve these objectives, the study employs several machine learning techniques, including regression models and ensemble methods such as Random Forest and XGBoost. These models are selected for their ability to process large, heterogeneous datasets and to capture complex nonlinear interactions among spatial, demographic, and infrastructural features. The machine learning analysis serves a dual purpose. First, it identifies the most profitable and strategically valuable locations for new charging stations by quantifying the influence of various factors—especially POI-related features—on station usage. This provides actionable guidance to charging point operators seeking to expand their networks in areas where user activity patterns naturally support higher demand. Second, the insights gained from the predictive analysis directly inform the formulation of the subsequent optimization model for optimal station placement. The learned relationships between POIs, traffic flows, and charging demand are thus embedded into the model's structure, ensuring that the optimization framework is grounded in empirically observed behavior rather than relying solely on theoretical assumptions.

5.1 Baseline Models

The first step of the analysis consisted in developing two baseline models, namely a basic Random Forest and a basic XGBoost classifier with no hyperparameter tuning. These models serve as reference points for evaluating the performance of the fine-tuned model developed later in the study. By relying on their default configurations, the baselines establish a neutral benchmark that reflects the predictive capability of these algorithms without any optimisation effort. Comparing the fine-tuned model against these baselines allows us to assess whether the methodological refinements introduced in this work lead to meaningful improvements or, alternatively, whether model performance deteriorates relative to simpler, untuned approaches.

5.1.1 Random Forest Model Development and Evaluation

As an initial step in our machine learning analysis, a baseline model was constructed using a `RandomForestRegressor`. This baseline serves as a foundational benchmark against which the performance of more advanced, fine-tuned models can be rigorously assessed. Its purpose is to evaluate the inherent predictive power of the selected features and to establish a neutral point of comparison prior to the introduction of any optimization techniques.

Model Setup

The baseline model was implemented using a `RandomForestRegressor` configured with 100 decision trees and a maximum depth of 5. These hyperparameters were intentionally kept simple to avoid overfitting while still providing sufficient representational capacity to capture the primary relationships within the dataset.

The dataset was first partitioned into a matrix of features, denoted by $\mathbf{X}_f$, and a target variable, denoted by y_f. The target variable represents the average number of daily charging sessions at each station and is computed as

$$y_f = \frac{S}{D - 30} \tag{5.1}$$

where S denotes the total number of charging sessions recorded at a station, and D represents the total number of days the station has been in operation. The first 30 days are excluded to account for the initial commissioning period. An 80/20 train–test split was applied, using a fixed random state to ensure reproducibility. This separation enables an unbiased evaluation of the model's generalization capability by measuring performance on data not seen during training.

5.1.2 XGBoost Model Development and Evaluation

For a robust initial assessment of predictive capabilities, a baseline model was developed using the XGBoost algorithm. XGBoost is renowned for its performance in regression tasks due to its efficient handling of various types of data and its ability to manage complex nonlinear relationships.

Model Configuration

The baseline XGBoost model was configured with the following parameters:

- **N_estimators:** Set to 1000 to allow the model to build up to 1000 trees, providing a substantial level of learning capacity and complexity.
- **Max_depth:** Limited to 3 to prevent the model from becoming overly complex and overfitting to the training data.
- **Learning Rate:** Set at 0.1 to control the weighting of new trees added to the model, helping to improve model stability and performance over iterations.
- **Objective:** Defined as 'reg:squarederror' to minimize the squared errors, which directly correlates with the RMSE evaluation metric used for performance assessment.
- **Early Stopping Rounds:** Configured to stop training if the validation score does not improve after 10 iterations, preventing unnecessary computations and potential overfitting.

The dataset was first split into training and testing sets, with 20% of the data reserved for testing. The training set was further split, allocating 25% as a validation set to monitor and tune the model's performance during training. The model was trained on the training set while periodically evaluated on the validation set using the RMSE metric. This strategy ensures that the model generalizes well to new data and identifies the best iteration for stopping before overfitting.

5.1.3 Model Evaluation

Following training, both models were evaluated on both the training and testing sets. Predictions were compared to observed values using several complementary performance metrics that quantify different aspects of predictive accuracy and reliability:

- **R^2 Score:** Measures the proportion of variance in the target variable explained by the model. Higher values indicate a stronger ability to capture the underlying structure of the data.
- **Root Mean Squared Error (RMSE):** Represents the average magnitude of prediction errors, giving greater weight to large deviations. Lower RMSE values reflect higher predictive precision.

- **Percentage Error:** Computed for both training and testing sets to express the average deviation between predicted and actual values as a percentage of the true values. This metric provides a more intuitive, scale-independent measure of model accuracy.

- **Feature Importance:** Derived from the Random Forest's built-in importance measure, this analysis quantifies the relative contribution of each feature to the prediction task, offering valuable insights into the explanatory relevance of key variables.

5.2 Fine-Tuned Models

Following the construction and evaluation of the baseline models, a second phase of the analysis focused on developing fine-tuned versions of both the Random Forest and XGBoost regressors. The goal of this phase was to enhance predictive performance by systematically optimizing model hyperparameters, thereby capturing complex interactions among features with greater accuracy. The improvements derived from these fine-tuned models are essential not only for achieving higher predictive reliability, but also for informing the subsequent optimization model for optimal charging station placement. Hyperparameter tuning was conducted using Optuna [2], a state-of-the-art framework for efficient and automated search over high-dimensional parameter spaces.

5.2.1 Random Forest Regression

The Random Forest regression model was selected for its robustness in handling non-linear relationships and its capacity to model interactions among a large number of explanatory variables. As an ensemble learning method, Random Forest constructs multiple decision trees and aggregates their outputs, offering high predictive accuracy and strong resistance to overfitting.

Hyperparameter Tuning

Building on the baseline RandomForestRegressor, an optimized model was developed through hyperparameter tuning using Optuna, a framework for automated hyperparameter optimization. Optuna formulates hyperparameter tuning as an optimization problem, where the objective is to maximize (or minimize) a predefined performance metric.

The method relies on a sequential search strategy based on Bayesian optimization, specifically the Tree-structured Parzen Estimator (TPE). At each iteration, Optuna selects a set of hyperparameters, trains the model, evaluates its performance, and uses the observed results to update a probabilistic model of the objective function. This allows the algorithm to efficiently explore the search space by prioritizing promising regions over less effective ones.

The search space for optimization included key parameters that influence the structure and behaviour of the forest:

- **Number of Estimators:** Varied between 50 and 1000 (in steps of 50) to balance model stability and computational efficiency.
- **Maximum Depth:** Explored between 3 and 10 to regulate the model's expressive power and mitigate overfitting.
- **Minimum Samples Split / Leaf:** Tuned to determine the minimum number of observations required to perform a split or to form a leaf node, affecting the granularity of the decision rules.
- **Max Features:** Examined strategies such as `sqrt`, `log2`, and `None` to assess how restricting the number of features available at each split affects model generalization.

An objective function was defined within the Optuna framework to instantiate a RandomForestRegressor with a given set of trial parameters and evaluate its performance using five-fold cross-validation. The mean R^2 score across folds served as the optimization criterion, ensuring that the selected model maximizes explained variance while maintaining strong generalization properties.

5.2.2 XGBoost Regression Model Development

This section presents the development of both the baseline and fine-tuned XGBoost models, designed to predict the usage intensity of EV charging stations. XGBoost, a gradient-boosted decision tree framework, is recognized for its efficiency, high predictive power, and ability to handle heterogeneous, tabular datasets with complex feature interactions.

The baseline model, implemented using default hyperparameters, provides a neutral reference point for assessing the impact of hyperparameter optimization. The fine-tuned model aims to surpass this benchmark by leveraging Optuna to explore a wide range of parameter configurations.

Hyperparameter Tuning with Optuna

To further improve predictive performance, an extensive hyperparameter search was conducted using Optuna. The optimization strategy focused on identifying the parameter combinations that minimize the Root Mean Squared Error (RMSE) obtained through cross-validation, thereby enhancing the model's accuracy and robustness.

The following key hyperparameters were included in the search space:

- **Max Depth:** Tested values between 3 and 6 to control tree complexity and reduce overfitting.

- **Subsample:** Varied between 0.6 and 0.95 to introduce randomness in row sampling and improve model generalization.
- **Colsample_bytree:** Tuned to regulate the proportion of features sampled for each tree, adding variability and reducing correlation among trees.
- **Lambda and Alpha:** L2 and L1 regularization weights were optimized to penalize overly complex models and enhance stability.

Each Optuna trial constructed an XGBoost model using the `DMatrix` data structure and evaluated its performance via five-fold cross-validation. The use of RMSE as the objective metric ensured that tuning emphasized minimizing errors in predicted charging demand, a crucial aspect for infrastructure planning applications.

5.2.3 Performance Evaluation

After hyperparameter tuning, the best-performing configurations identified by Optuna were used to train the final Random Forest and XGBoost models. These tuned models were then evaluated on both the training and testing sets to assess their predictive accuracy, generalizability, and robustness. A comprehensive suite of performance metrics was employed:

- **R^2 Score:** The coefficient of determination was computed for both training and testing sets to quantify the proportion of variance in the average number of charging sessions that the models were able to explain.
- **RMSE:** The Root Mean Squared Error provided an absolute measure of prediction error, offering a direct and interpretable indication of model precision.
- **Percentage Error:** This metric measures the average percentage deviation between predicted and actual values, providing an intuitive, scale-independent indicator of practical performance.
- **Feature Importance:** Feature importance scores were examined to identify the variables that contributed most significantly to the predictive task, offering insights into the underlying factors influencing station usage.

5.3 Results from the Deviation-Flow-Based Grouping Approach

The final phase of the analysis focuses on the results obtained when applying the fine-tuned models to the grouped dataset derived from the deviation-flow–based methodology. This grouping strategy—developed as an alternative to the POI-based classification—organizes node pairs according to the deviation flows captured between origin–destination combinations. By structuring the dataset to reflect realistic movement patterns across the region, this approach provides a more behaviourally informed representation of demand, allowing the models to account for traffic-driven spatial dynamics that influence the attractiveness and usage of EV charging stations.

To rigorously evaluate the stability and robustness of the models under this grouping framework, the dataset was repeatedly partitioned into training and testing subsets using twelve distinct random states: $[0, 1, 15, 16, 41, 42, 52, 53, 61, 62, 100, 101]$.
Employing multiple random states ensures that model performance is not an artifact of a particular data split. Instead, it enables a more comprehensive assessment of generalizability by observing how predictive accuracy varies across different train–test configurations. This approach is particularly relevant given the spatial and behavioural heterogeneity embedded within the deviation-flow–based groups.

Table 5.1: Summary of Model Performance for Deviation-Flow Capturing Approach Across Different Random States

RS	Model	Phase	R^2 Train	R^2 Test	% Error Train	% Error Test
0	RF	Baseline	0.7578	0.5269	253.19%	204.25%
	RF	Optimized	0.8842	0.5564	115.14%	126.36%
	XGB	Baseline	0.7805	0.5645	193.14%	158.81%
	XGB	Optimized	0.9775	0.4936	80.13%	147.42%
1	RF	Baseline	0.7648	0.5486	257.83%	243.02%
	RF	Optimized	0.9023	0.6065	126.81%	156.00%
	XGB	Baseline	0.8923	0.6238	150.04%	163.24%
	XGB	Optimized	0.9738	0.6382	76.58%	161.00%
15	RF	Baseline	0.7661	0.6258	235.61%	294.70%
	RF	Optimized	0.8312	0.6888	112.67%	190.45%
	XGB	Baseline	0.6931	0.5372	358.48%	426.61%
	XGB	Optimized	0.8700	0.5822	135.72%	199.26%
16	RF	Baseline	0.7582	0.6095	246.93%	230.60%
	RF	Optimized	0.8467	0.5409	120.68%	148.68%
	XGB	Baseline	0.8388	0.5151	172.50%	188.84%
	XGB	Optimized	0.9308	0.6568	123.79%	168.14%
41	RF	Baseline	0.7479	0.6347	238.41%	193.51%
	RF	Optimized	0.9118	0.6999	112.83%	141.15%
	XGB	Baseline	0.8023	0.5582	204.67%	169.25%
	XGB	Optimized	0.9349	0.7020	110.52%	169.37%
42	RF	Baseline	0.7509	0.6592	240.81%	232.77%
	RF	Optimized	0.7813	0.6548	126.68%	158.44%
	XGB	Baseline	0.8685	0.5948	156.90%	198.40%

Table 5.1 – continued from previous page

RS	Model	Phase	R^2 Train	R^2 Test	% Error Train	% Error Test
	XGB	Optimized	0.9492	0.6964	97.86%	178.23%
52	RF	Baseline	0.7756	0.4533	241.62%	248.00%
	RF	Optimized	0.8807	0.4854	147.98%	170.16%
	XGB	Baseline	0.7845	0.5722	195.88%	200.29%
	XGB	Optimized	0.8764	0.5852	152.38%	185.77%
53	RF	Baseline	0.7498	0.5821	223.57%	274.95%
	RF	Optimized	0.9150	0.6067	99.27%	228.60%
	XGB	Baseline	0.8470	0.5582	158.72%	214.70%
	XGB	Optimized	0.9789	0.5009	68.54%	276.95%
61	RF	Baseline	0.7759	0.4712	221.54%	262.94%
	RF	Optimized	0.9170	0.5290	108.92%	207.18%
	XGB	Baseline	0.7517	0.4783	263.18%	328.14%
	XGB	Optimized	0.8986	0.5568	152.41%	221.25%
62	RF	Baseline	0.7439	0.6430	243.25%	230.57%
	RF	Optimized	0.7271	0.6856	133.16%	153.02%
	XGB	Baseline	0.7737	0.6394	202.58%	190.75%
	XGB	Optimized	0.9419	0.6476	121.94%	159.34%
100	RF	Baseline	0.7570	0.5487	238.22%	315.43%
	RF	Optimized	0.9033	0.5875	118.85%	189.68%
	XGB	Baseline	0.8652	0.5073	144.79%	211.36%
	XGB	Optimized	0.8909	0.5809	124.91%	200.62%
101	RF	Baseline	0.7708	0.5729	228.80%	236.38%
	RF	Optimized	0.9103	0.6214	104.98%	172.80%
	XGB	Baseline	0.8457	0.6184	189.88%	210.88%
	XGB	Optimized	0.9869	0.5655	61.77%	196.79%

5.3.1 Discussion on Model Stability and Performance Variability

As shown in table 5.1, a notable finding from the evaluation of the predictive models is the substantial variability in R^2 values obtained across the twelve different random train–test splits, with test-set performance ranging from 0.4854 to 0.7020. Such a broad interval raises concerns regarding the stability and reliability of the models, suggesting a marked dependence on the specific composition of the training and testing subsets.

The observed fluctuations indicate that, although the models are capable of achieving relatively high predictive accuracy under certain data configurations, their performance deteriorates considerably under others. This instability is indicative of a degree of sensitivity to the underlying data distribution and may signal potential overfitting, whereby the models inadvertently capture noise or artifacts present in particular splits rather than learning generalizable patterns.

Several factors likely contribute to this performance variability:

- **Weak Feature Relationships:** Pearson and Kendall correlation analyses revealed only weak associations between Points of Interest (POIs) and charging station usage. This suggests that the explanatory power of the initial feature set is limited, restricting the models' ability to consistently capture the determinants of charging demand.
- **Issues Identified in the Initial Feature Engineering Process:**
  - **Overly Broad POI Grouping:** The initial classification aggregated POIs into only six general categories. This lack of granularity likely obscured meaningful distinctions between POI types, reducing the ability of the model to detect more nuanced behavioural patterns that influence EV charging activities.
  - **Inappropriate Radius Selection:** The initial use of a 10-minute radius—while intuitively reasonable for larger urban areas—proved excessively broad for smaller municipalities. This led to the inclusion of POIs that were not realistically accessible or relevant to actual charging behaviour, thereby introducing noise into the feature set.

Overall, the combination of weak feature–target relationships, inadequate spatial granularity, and heterogeneous urban contexts contributed to the instability of the models. These insights highlight the need for a more refined approach to feature engineering—particularly in relation to POI categorization and spatial relevance—to enhance model robustness and improve the reliability of subsequent predictive analyses.

5.4 Results from the Node-Serving Model Approach

In this section, we present the results obtained using the *node-serving model*, an alternative grouping strategy grounded in the principles of deviation-flow modelling. Unlike the initial POI-based approach, which relied primarily on the spatial distribution and categorization of nearby amenities, the node-serving model leverages the structure of the transportation network itself. Its objective is to identify how pairs of nodes (cities or municipalities) interact through observed or estimated deviation flows, and to group charging stations according to the demand they are most likely to serve along these flows. This method provides a behaviourally informed perspective on charging demand, reflecting mobility patterns rather than solely the static characteristics of surrounding urban environments.

The theoretical motivation for this approach stems from the need to capture realistic usage patterns of EV charging infrastructure. Since charging demand is intrinsically linked to travel behaviour, grouping stations by their role in serving deviation flows enables the predictive models to incorporate origin–destination dynamics, roadway accessibility, and spatial load distribution. This is particularly relevant in regions where functional mobility connections play a stronger role in determining station usage than local POI density alone.

The retrieval of the mobility-related information required for this grouping, as well as the construction of the final set of features, was discussed in detail in Chapter 2. The present section builds upon that foundation by applying the node-serving grouping methodology to the refined dataset and evaluating its impact on the performance and stability of the predictive models.

To ensure methodological consistency and allow a direct comparison with the previous experimental results, the models were trained using the same computational pipeline employed in earlier analyses. All steps—including data preprocessing, feature engineering, model training, and evaluation—were kept identical across experiments, with the only difference being the grouping logic that underlies this new dataset. This controlled setup enables an objective assessment of whether the node-serving model leads to more stable or accurate predictions than the deviation-flow–based and POI-based approaches.

5.4.1 Discussion of Results and Model Behaviour

The application of the node-serving grouping methodology did not yield substantial improvements in predictive performance, as shown in Table 5.2. Despite the shift from a POI-centred representation to a mobility-driven approach, the R^2 values of the fine-tuned models remained within a range comparable to previous experiments, varying between 0.4414 and 0.7203 across the different random train–test splits. This interval mirrors the variability observed with the original grouping methods, indicating that the reformulation of the dataset did not significantly enhance the models' ability to explain the variance in charging station usage.

Table 5.2: Summary of Model Performance for Node-serving Approach Across Different Random States

RS	Model	Phase	R^2 Train	R^2 Test	% Error Train	% Error Test
0	RF	Baseline	0.7532	0.5522	242.64%	192.35%
	RF	Optimized	0.9108	0.5965	121.36%	119.16%
	XGB	Baseline	0.7932	0.6063	175.11%	147.56%
	XGB	Optimized	0.8536	0.5771	149.24%	134.33%
1	RF	Baseline	0.7651	0.4812	247.15%	230.56%
	RF	Optimized	0.8842	0.5970	122.95%	128.99%
	XGB	Baseline	0.7849	0.6388	246.37%	212.77%
	XGB	Optimized	0.9697	0.6240	82.40%	159.10%
15	RF	Baseline	0.7608	0.6102	230.86%	302.49%
	RF	Optimized	0.8549	0.6741	113.36%	188.70%
	XGB	Baseline	0.6110	0.5686	416.00%	529.33%
	XGB	Optimized	0.9139	0.5676	115.56%	243.12%

Table 5.2 – continued from previous page

RS	Model	Phase	R^2 Train	R^2 Test	% Error Train	% Error Test
16	RF	Baseline	0.7519	0.6321	242.93%	230.93%
	RF	Optimized	0.8729	0.6974	121.41%	140.19%
	XGB	Baseline	0.8382	0.5644	170.50%	173.85%
	XGB	Optimized	0.9186	0.6785	120.40%	178.45%
41	RF	Baseline	0.7456	0.5967	234.53%	194.24%
	RF	Optimized	0.7927	0.6334	123.94%	139.60%
	XGB	Baseline	0.8632	0.6213	155.53%	150.37%
	XGB	Optimized	0.9285	0.6764	118.65%	167.24%
42	RF	Baseline	0.7494	0.6651	238.85%	226.28%
	RF	Optimized	0.8175	0.6990	122.58%	153.72%
	XGB	Baseline	0.8006	0.6637	185.27%	185.80%
	XGB	Optimized	0.8512	0.6570	178.02%	170.37%
52	RF	Baseline	0.7715	0.3957	235.66%	246.22%
	RF	Optimized	0.8821	0.5191	117.54%	152.73%
	XGB	Baseline	0.7410	0.5790	222.94%	207.38%
	XGB	Optimized	0.9517	0.6226	102.56%	165.22%
53	RF	Baseline	0.7450	0.5735	214.97%	266.97%
	RF	Optimized	0.9124	0.6105	98.67%	226.51%
	XGB	Baseline	0.8332	0.5748	153.23%	238.21%
	XGB	Optimized	0.9604	0.5310	78.96%	262.74%
61	RF	Baseline	0.7709	0.4610	219.23%	261.94%
	RF	Optimized	0.8807	0.5224	108.95%	203.02%
	XGB	Baseline	0.6864	0.4640	328.14%	381.74%
	XGB	Optimized	0.8632	0.4414	144.78%	203.47%
62	RF	Baseline	0.7428	0.6552	237.42%	231.83%
	RF	Optimized	0.6936	0.6816	141.77%	154.39%
	XGB	Baseline	0.8006	0.6682	178.81%	190.20%
	XGB	Optimized	0.9731	0.7203	78.69%	148.28%
100	RF	Baseline	0.7472	0.5736	227.92%	290.80%
	RF	Optimized	0.8969	0.5959	121.46%	184.27%
	XGB	Baseline	0.8331	0.5208	153.28%	213.02%
	XGB	Optimized	0.9634	0.5672	83.67%	225.47%
101	RF	Baseline	0.7569	0.5993	221.31%	227.97%
	RF	Optimized	0.8452	0.6533	108.38%	171.85%
	XGB	Baseline	0.8738	0.6094	162.40%	207.93%
	XGB	Optimized	0.9728	0.6307	72.68%	185.98%

Several hypotheses may account for the limited impact of this new approach:

- **Insufficient Signal in Deviation-Flow Features:** While the node-serving strategy incorporates behavioural information derived from origin–destination flows, these flows may not be sufficiently predictive of EV charging demand at the municipal scale. If deviation flows are relatively homogeneous or weakly correlated with actual charging behaviour, the resulting features might contribute little additional explanatory power.

- **Persistence of Weak Relationships Between Features and Target:** As discussed in Chapter 2, both Pearson and Kendall analyses revealed weak associations between POIs and charging demand. Even after reorganizing the dataset through deviation-flow logic, the underlying features may still lack the strong, direct relationships needed to support a highly predictive model. This suggests that the core predictors available—whether POI-based or flow-based—may not capture the dominant drivers of EV charging behaviour.

- **Limitations of the Spatial Resolution:** The node-serving approach operates at the level of municipalities or aggregated nodes. It is plausible that important micro-scale dynamics—such as precise station accessibility, micro-mobility patterns, or neighbourhood-level socio-economic factors—are lost when data are aggregated to this scale. As a result, the model may not fully capture local heterogeneity that strongly influences usage.

- **High Variability in Charging Demand:** EV charging demand is inherently volatile and driven by factors that are difficult to encode in static spatial features alone (e.g., temporal patterns, driver routines, weather, energy prices). This intrinsic variability may limit the upper bound of achievable predictive performance regardless of the grouping approach.

- **Residual Sensitivity to Train–Test Splits:** The persistence of large fluctuations in R^2 across random states suggests that the model remains sensitive to the specific partitioning of the dataset. This is consistent with a scenario in which the amount of signal in the features is limited relative to the noise, leading to unstable performance even when the grouping logic is modified.

Overall, the results indicate that transitioning to a node-serving perspective—while conceptually grounded in mobility behaviour—does not resolve the instability or performance limitations observed in earlier models. This suggests that future improvements may require either (i) the incorporation of additional, more informative features (e.g., temporal usage patterns, electricity market data, user behaviour data), or (ii) the adoption of modelling strategies capable of capturing more complex spatial and relational structures, such as graph-based or deep learning approaches.

5.5 Binary Classifier

Given the instability and limited explanatory power observed in the regression models, a final modelling strategy was developed based on a binary classification framework. The aim of this approach is not to predict the exact number of charging sessions, but rather to classify each station as *profitable* or *non-profitable* according to a threshold grounded in empirical evidence. This formulation aligns more naturally with the practical decision-making problem faced by charging point operators, who often must determine whether a location justifies investment rather than requiring a precise demand estimate.

5.5.1 Definition of the Profitability Threshold

To construct a binary target variable, it was necessary to define a profitability threshold α capable of distinguishing between high-performing and low-performing stations. Drawing on the feature importance analysis from the fine-tuned XGBoost model—where the *number of connectors* and *connector power* emerged as two of the most influential predictors—profitability was conceptualized in terms of usage *relative to the installed power capacity* of each station.

The following procedure was adopted:

1. **Compute an empirical α for each station.**
   - Derive the total installed power:
   $$\text{total\_power} = n\_cp \times \text{conn\_power}.$$
   - Assign each station to a power range: *50–99 kW*, *100–149 kW*, or *150+ kW*.
   - Compute
   $$\alpha = \frac{\text{avg\_tot\_sessions}}{\text{total\_power}},$$
   which measures how intensively each kilowatt of capacity is utilized.

2. **Summarize α within each power range.** For each power segment, both the mean and median of the empirical α values were computed:
   - **50–99 kW:** median = 0.0089, mean = 0.0129
   - **100–149 kW:** median = 0.0108, mean = 0.0133
   - **150+ kW:** median = 0.0057, mean = 0.0071

3. **Define binary labels.** For each station, and for each chosen α statistic (mean or median), a binary label was assigned:
$$\text{label}_\alpha = \begin{cases} 1, & \text{if avg\_tot\_sessions} \geq \alpha_{\text{stat}} \times \text{total\_power}, \\ 0, & \text{otherwise.} \end{cases}$$

Both the mean- and median-based thresholds were retained for experimentation, while higher percentiles (e.g., 70th, 80th, 90th) were discarded due to their negative impact on accuracy and excessive inflation of precision, which generated overly conservative models with limited operational usefulness.

5.5.2 Experimental Setup

For each threshold definition (mean- and median-based α), a series of binary classification experiments were conducted using a fine-tuned XGBoost classifier. The evaluation followed the same multi–random-state protocol adopted in the regression experiments to ensure robustness. Performance was assessed through:

- **Accuracy:** Measures the proportion of correctly classified instances out of the total number of predictions. It gives an overall indication of model correctness but may be misleading in imbalanced datasets.
- **Precision (class 1 and class 0):** Precision quantifies how many of the instances predicted to belong to a given class are actually correct. For class 1 (profitable stations), it expresses the reliability of positive predictions; for class 0, it reflects the reliability of negative predictions.
- **Recall (class 1 and class 0):** Recall measures the proportion of actual instances of a class that the model correctly identifies. For class 1, it indicates how effectively the model retrieves profitable stations; for class 0, it measures how well it detects non-profitable stations.
- **F1-score:** The harmonic mean of precision and recall. It provides a balanced measure that accounts for both false positives and false negatives, making it especially useful when class distributions are not perfectly balanced.
- **ROC–AUC:** The Area Under the Receiver Operating Characteristic Curve. It evaluates the model's ability to discriminate between the two classes across all possible classification thresholds. Higher values indicate stronger separability between profitable and non-profitable stations.
- **Mean Confidence for Class 1 Predictions:** Represents the average predicted probability assigned to the profitable class. It reflects how confident the classifier is when identifying stations as profitable.
- **Log-loss and Error Rate:** Log-loss penalizes incorrect predictions based on their probability estimates, rewarding confident and correct predictions while heavily penalizing confident but incorrect ones. The error rate complements accuracy by measuring the proportion of incorrect predictions.

This comprehensive suite of metrics enabled an evaluation not only of predictive correctness, but also of the confidence, stability, and discriminative power of the classifier.

Table 5.3: Performance of the XGBoost Binary Classifier under Mean- and Median-Based α Thresholds

(A) Mean-Based α Threshold

Label	Accuracy	Prec.1	Prec.0	Rec.1	Rec.0	F1	ROC-AUC	Conf.1
logloss	0.666	0.563	0.733	0.575	0.723	0.568	0.713	0.749
error	0.684	0.682	0.687	0.693	0.675	0.687	0.747	0.683
auc	0.686	0.683	0.689	0.695	0.675	0.689	0.748	0.681

(B) Median-Based α Threshold

Label	Accuracy	Prec.1	Prec.0	Rec.1	Rec.0	F1	ROC-AUC	Conf.1
logloss	0.683	0.680	0.687	0.695	0.671	0.687	0.744	0.679
error	0.671	0.569	0.736	0.577	0.729	0.573	0.719	0.751
auc	0.669	0.567	0.733	0.570	0.730	0.568	0.714	0.760

5.5.3 Classification Results

The results for the two thresholding strategies are summarised in table 5.3 below.

Although the quantitative performance metrics suggest that the combinations *mean + AUC* or *mean + error* perform similarly to *median + log-loss*, a closer inspection of the feature importance rankings reveals important qualitative differences that favour the median-based threshold. In particular, when analyzing the POI-related features that consistently appear among the top 10 most important predictors across more than half of the random states, the *median + log-loss* configuration shows a more coherent and interpretable pattern. For this model, the most influential POIs include *fast food*, *mall*, *supermarket*, and *sustenance_long*—a group that aligns well with what is typically observed when inspecting the 100 most profitable stations manually. These POI categories capture locations where users naturally spend a significant amount of time, making them intuitively compatible with EV charging dwell times.

By contrast, the configurations based on the mean threshold show a less consistent and less behaviourally meaningful set of important features. The *mean + error* model identifies only three POIs—*amenity_others*, *mall*, and *wholesale*—while the *mean + AUC* model retains only *mall* and *wholesale*. The reduced diversity and lower interpretability of these POI-driven predictors suggest that the mean-based classifications may fail to capture the subtler behavioural signals underlying station profitability. This reinforces the idea that the median threshold provides a more stable and realistic characterization of what constitutes a profitable station within the dataset.

A particularly notable finding is not simply which POIs appear as important, but which do not. Categories such as *bar*, *house*, *office*, or *sport facilities*—which might intuitively seem associated with routine mobility—do not emerge as strong predictors of profitability. This suggests that stations in areas dominated by these POIs, despite potentially high footfall or activity levels, do not attract the sustained dwell times or charging behaviours

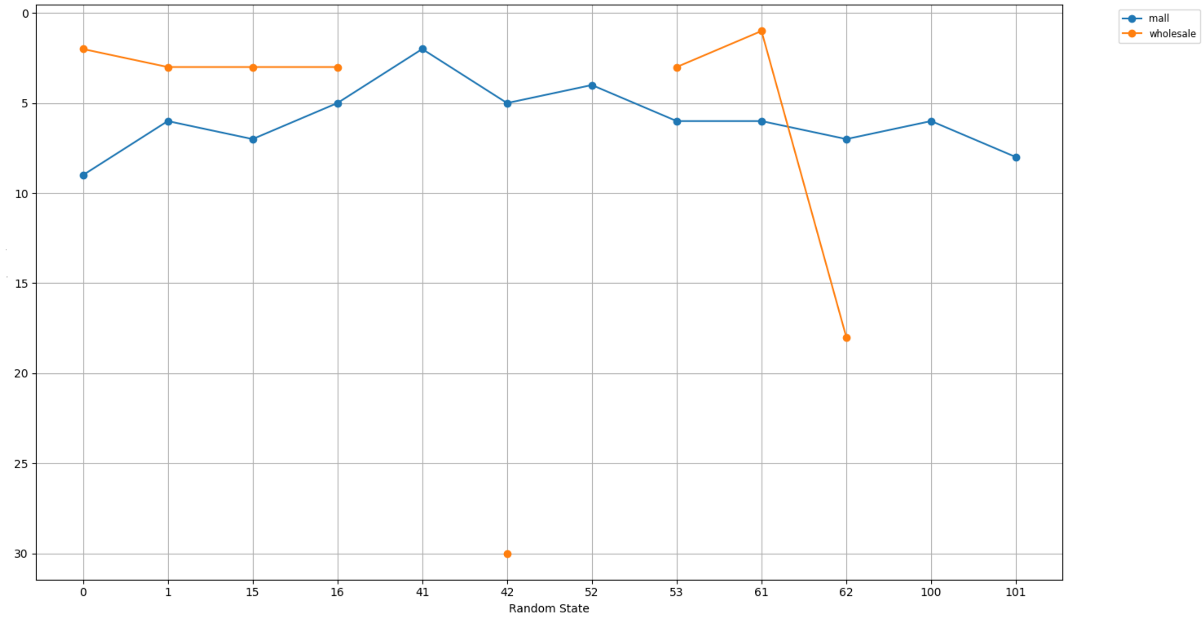

Figure 5.1: Ranking of the selected features across the different random states for mean + AUC

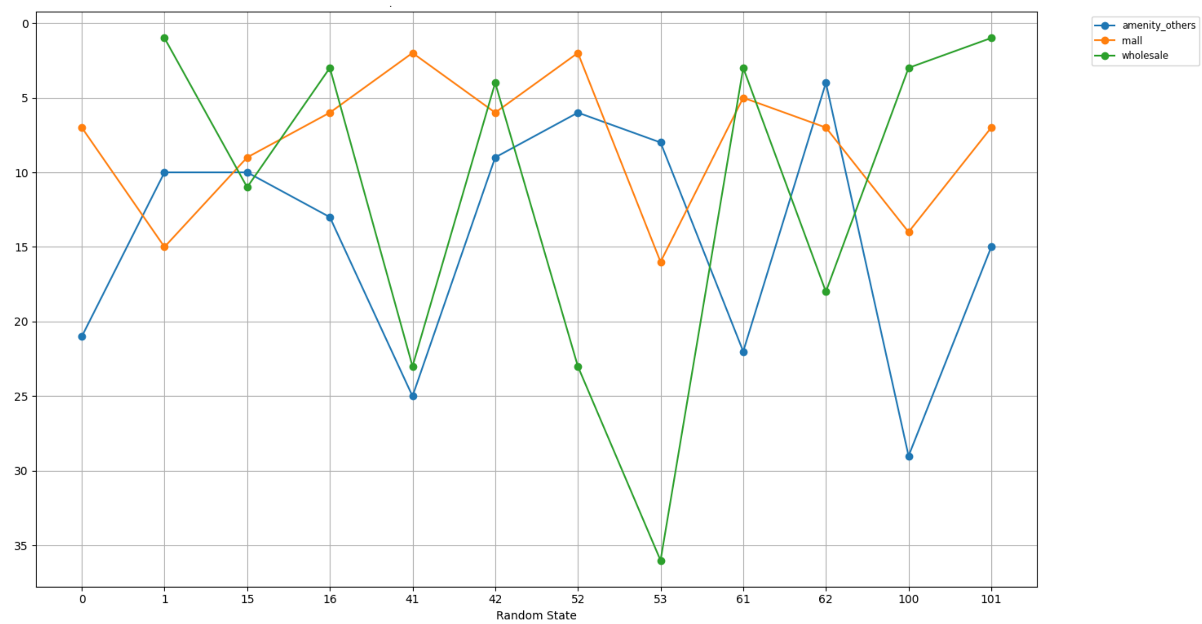

Figure 5.2: Ranking of the selected features across the different random states for mean + error

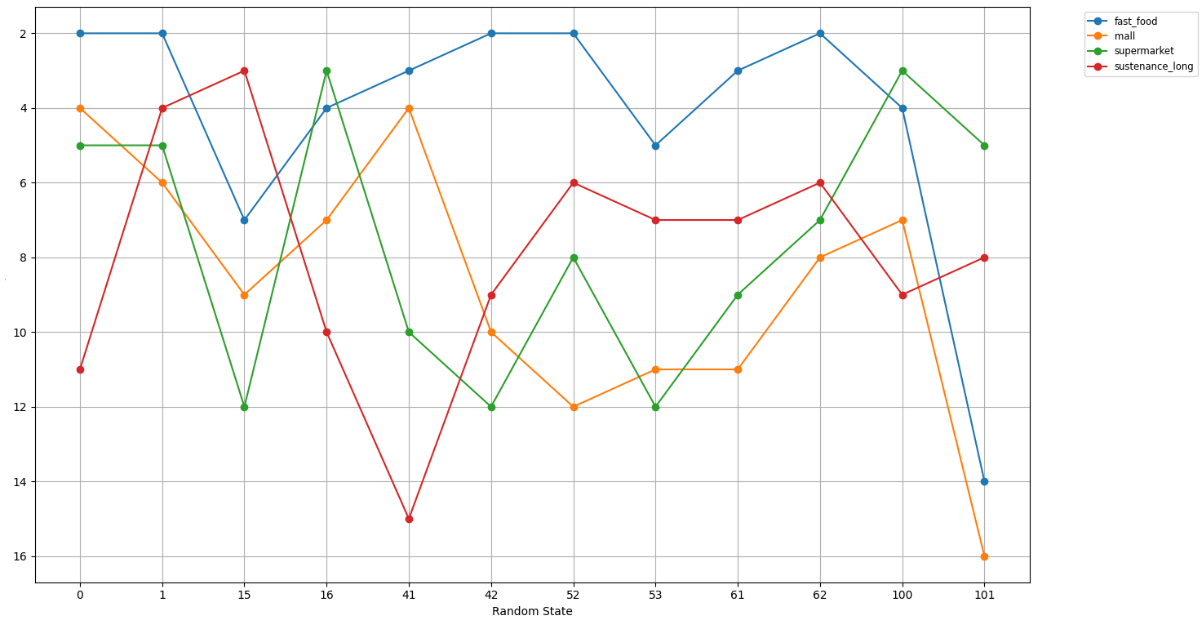

Figure 5.3: Ranking of the selected features across the different random states for median + log-loss

Table 5.4: Performance Metrics for the Median-Based α Threshold Across Random States

Random State	Accuracy	Precision1	Precision0	Recall1	Recall0	F1-score	ROC-AUC	Conf.1
0	0.658	0.661	0.654	0.650	0.665	0.655	0.734	0.693
1	0.698	0.690	0.707	0.722	0.674	0.705	0.760	0.666
15	0.691	0.687	0.696	0.705	0.678	0.696	0.740	0.693
16	0.687	0.693	0.682	0.675	0.699	0.684	0.756	0.661
41	0.704	0.703	0.705	0.709	0.699	0.706	0.766	0.696
42	0.666	0.654	0.681	0.709	0.623	0.680	0.705	0.709
52	0.689	0.705	0.676	0.654	0.725	0.678	0.748	0.660
53	0.660	0.652	0.668	0.688	0.631	0.669	0.730	0.672
61	0.660	0.664	0.656	0.650	0.669	0.657	0.726	0.694
62	0.685	0.675	0.697	0.717	0.653	0.695	0.745	0.661
100	0.712	0.704	0.721	0.734	0.691	0.719	0.767	0.669
101	0.681	0.668	0.696	0.722	0.640	0.694	0.746	0.675
Mean	0.6826	0.6797	0.6866	0.6946	0.6706	0.6865	0.7436	0.6791

Table 5.5: Feature Importance Ranks Across Random States for the Binary Classifier

Feature	**0**	**1**	**15**	**16**	**41**	**42**	**52**	**53**	**61**	**62**	**100**	**101**
RANDOM_INT	23	24	22	18	30	21	31	26	24	20	26	25
amenity_facilities	17	29	27	12	17	17	24	21	13	13	21	11
amenity_others	10	34	28	23	8	32	7	15	31	24	16	35
block_traffic	7	7	2	2	6	4	5	3	6	4	2	3
count_in_site_DC	9	10	5	9	11	11	9	14	8	5	6	12
department_store	15	35	-	30	34	28	18	31	22	25	24	32
distance_to_primary	13	16	17	15	18	16	20	10	5	15	10	18
education	22	30	19	28	19	34	27	27	27	34	20	13
electric_cars	16	27	14	14	29	14	28	19	18	22	23	31
entertainment	33	28	33	29	13	7	33	22	10	17	33	15
fast_food	2	2	7	4	3	2	2	5	3	2	4	14
financial	29	15	23	20	20	20	25	32	4	31	35	7
fuel	30	14	30	26	14	3	14	23	28	30	34	24
furniture	8	3	11	-	25	30	26	4	35	18	29	4
healthcare	25	32	4	24	24	24	13	17	17	33	15	27
house	28	26	24	21	23	23	22	29	26	21	25	22
leisure	31	21	20	25	16	26	23	6	12	16	19	9
mall	4	6	9	7	4	10	12	11	11	8	7	16
motorway_flag	3	11	31	34	31	29	3	-	29	32	5	10
office	32	33	32	19	35	6	32	33	14	23	32	33
on_the_go	36	22	34	32	9	35	-	34	9	18	21	-
operation_days	1	1	1	1	1	1	1	1	1	1	1	1
parking	26	31	21	13	28	27	29	13	30	28	30	17
population_comune	6	8	6	5	2	8	4	8	2	3	8	2
population_prov	19	20	16	17	7	18	15	25	19	11	13	29
population_reg	14	9	10	8	5	15	16	9	16	12	11	23
public_service	34	17	25	33	26	25	35	24	33	26	28	30
shop	20	13	15	6	12	5	10	16	15	10	12	6
spiritual	21	23	29	31	33	33	34	30	32	35	31	34
sports	18	19	13	27	21	19	19	28	20	19	27	26
supermarket	5	5	12	3	10	12	8	12	9	7	3	5
sustenance_long	11	4	3	10	15	9	6	7	7	6	9	8
sustenance_short	27	25	26	22	36	22	21	20	25	27	14	28
tourism	12	18	8	11	27	13	11	18	23	14	22	19
transportation	24	12	18	16	32	31	17	2	21	29	17	20
wholesale	35	-	-	-	22	-	30	-	-	-	-	-

required to become consistently profitable. Instead, the results point toward commercial complexes and food-related activities as more significant attractors, likely because they align better with typical charging durations and user behaviour during charging stops.

Beyond POI-related variables, several structural and contextual features consistently emerged as important predictors of station profitability across the different random states. These features do not describe the commercial environment surrounding a station, but rather its operational characteristics, spatial context, and relationship with the transportation network. Their prominence in the feature-importance rankings provides valuable insight into the fundamental factors driving EV charging behaviour.

Operation Days. Across all random states, *operation_ days* appeared as the most important feature. This result is intuitive: stations that have been active for a longer period naturally accumulate more charging sessions, but they also tend to be better integrated into user routines and navigation apps, making them more visible and more frequently chosen by drivers. The dominance of this feature highlights the temporal dimension of infrastructure maturity—stations require time to build a user base, meaning that longevity itself is a strong proxy for demand.

Count in Site (DC). The number of charging points available at a site reflects its capacity to simultaneously accommodate multiple vehicles. Stations with a larger number of connectors (particularly DC fast chargers) are better positioned to serve peak demand and reduce waiting times, making them more attractive to drivers. The consistent importance of *count_in_site_DC* suggests that capacity plays a central role in shaping usage, especially in areas where demand surges during commuting hours or along mobility corridors.

Block Traffic. The *block_traffic* feature—reflecting the load and flow intensity of the road segment where the station is located—also appears frequently among the top predictors. This aligns with the idea that EV charging behaviour is inherently tied to mobility patterns: stations situated along high-traffic blocks are exposed to a larger potential user base and therefore tend to exhibit higher utilization. This finding strengthens the broader conclusion that mobility-driven features, not only POIs, are essential for understanding demand.

Parking Availability. The *parking* feature captures the availability of parking spaces near the station. Its recurring importance indicates that accessible, convenient parking remains a key determinant of station attractiveness. Users prefer locations where they can easily stop, connect their vehicle, and leave it charging without logistical constraints. Limited or inconvenient parking can deter charging, even in otherwise favourable contexts.

Population of the Municipality. The feature *population_comune* also ranked consistently among the most important predictors. This indicates that local population density or municipal scale still plays a role in shaping charging demand, even after controlling for POIs and mobility flows. While POIs did not show strong linear correlations with usage, the broader demographic scale of a municipality likely captures diffuse socio-economic and behavioural factors tied to EV ownership and mobility demand.

Taken together, these features highlight that the drivers of station profitability extend beyond the immediate commercial surroundings. Operational maturity, infrastructural capacity, accessibility, and demographic context emerge as equally—if not more—relevant factors.

Finally, a critical qualitative observation further supports this interpretation: among the stations in the 90th percentile of profitability, virtually all are located directly on major highways or in close proximity to them. This confirms that high profitability is strongly associated with long-distance mobility corridors, where charging stops are integrated into travel itineraries and where the opportunity cost of charging aligns naturally with rest or consumption activities.

5.5.4 Range-Specific Classification Models and Empirical Threshold Definition

To enhance the strategic value of the analysis and provide actionable guidance for infrastructure planning, the final phase of the study involved developing dedicated binary classification models for different categories of charging stations. Since stations with different power levels (50–99 kW, 100–149 kW, and 150+ kW) serve distinct operational roles and user needs, a single global profitability threshold would be inadequate. Instead, tailoring the predictive model to each power range allows the decision process to reflect the intrinsic characteristics and performance expectations of stations with similar technological configurations.

To construct these range-specific binary labels, we followed a systematic approach for deriving an empirical profitability threshold, denoted as α, for each dataset. The threshold captures the ratio between charging demand and installed capacity, thus reflecting how intensively each kilowatt of power is utilized. For every station, we computed:

$$\alpha_{\text{emp}} = \frac{\text{avg\_tot\_sessions}}{\text{total\_power}}, \qquad \text{where} \quad \text{total\_power} = n_{cp} \times \text{conn\_power}.$$

This metric normalizes usage by capacity, allowing stations of different sizes to be compared in a meaningful and consistent way. For each power range, we derived the distribution of empirical α values and extracted key statistics—the mean, median, 25th, 75th, and 90th percentiles. These statistics were then used as candidate thresholds to explore the effect of different profitability definitions on the classification performance.

The resulting values show clear systematic differences between power ranges: lower-power stations tend to exhibit higher α values (i.e., higher usage per installed kilowatt), while high-power stations show lower α values, consistent with their role in serving rapid, transient charging needs rather than maximizing occupancy. These patterns confirm that profitability is not absolute but must be interpreted relative to the technical characteristics of each station segment.

The empirical thresholds obtained for each power category are summarized in Table 5.6.

Table 5.6: Empirical α thresholds (sessions per kW of installed power) for each station power range

Power Range (kW)	**25th pct.**	**Median**	**Mean**	**75th pct.**	**90th pct.**
50–99 kW	0.0035	0.0089	0.0129	0.0183	0.0304
100–149 kW	0.0055	0.0108	0.0133	0.0179	0.0262
150+ kW	0.0028	0.0057	0.0071	0.0092	0.0142

To evaluate the performance of the binary classifier under different profitability thresholds, we computed the mean metrics obtained across all random states for each empirical α value, metric type (AUC, error, log-loss), and power range. The results, reported in Table 5.7, show a consistent pattern: thresholds based on the **median** values of α offer the best trade-off between accuracy, precision, recall, and ROC–AUC across all three power ranges. Higher percentile thresholds (such as the 75th or 90th percentile) lead to inflated precision but significantly reduced recall and overall performance, making them less suitable for operational decision-making. Therefore, median-based thresholds were selected as the most robust and balanced definitions of station profitability.

The analysis of Points of Interest (POI) feature importance across all combinations of power range and evaluation metric reveals several consistent patterns. These patterns provide insight into how user behaviour and local context influence station profitability at different charging capacities.

1. Power range 50–99 kW

Across AUC, Error, and Log-loss metrics, the following POIs consistently appear among the top ten most relevant features: **distance_to_primary**, **motorway_flag**, **fast_food**, and **mall**.

These POIs collectively indicate that medium-power charging stations tend to be most profitable when they are:

- located near major roads,
- situated close to fast-food chains (typical intermediate dwell-time activities),
- positioned near commercial centres such as malls.

This aligns with expected user behaviour: drivers using mid-power chargers typically plan stops of intermediate duration, which match well with fast-food establishments and shopping malls. Proximity to motorways further suggests that this power range partially serves transit-oriented charging demand.

Additional overlaps include:

- **sustenance_long**, shared between Log-loss and Error,
- **transportation**, shared between AUC and Error.

These features highlight that both long-dwell food options and multimodal transport hubs contribute meaningfully to station profitability.

2. Power range 100–149 kW

For this intermediate-high power category, the following POIs appear across all three metrics: **house** and **tourism**.

These results indicate that stations in this power category are most used in:

- residential environments, where EV ownership may be higher and home-charging constraints can lead to public charging demand,

Table 5.7: Mean performance metrics across empirical α percentile thresholds, power ranges, and evaluation metrics.

Power	Metric	Percentile	Acc.	Prec.1	Prec.0	Rec.1	Rec.0	F1	ROC-AUC	Conf.1
50–99 kW	AUC	25th	0.7124	0.8105	0.4325	0.8033	0.4428	0.8068	0.7214	0.8588
		Median	0.6918	0.6879	0.6971	0.7072	0.6764	0.6969	0.7520	0.6918
		Mean	0.6786	0.5618	0.7556	0.5997	0.7248	0.5798	0.7164	0.7576
		75th	0.7337	0.4733	0.8310	0.5117	0.8086	0.4912	0.7413	0.7774
		90th	0.8505	0.2486	0.9170	0.2501	0.9173	0.2482	0.7062	0.7579
	Error	25th	0.7127	0.8083	0.4302	0.8076	0.4312	0.8080	0.7148	0.8560
		Median	0.6896	0.6862	0.6942	0.7042	0.6751	0.6948	0.7537	0.6873
		Mean	0.6715	0.5538	0.7478	0.5838	0.7228	0.5679	0.7143	0.7543
		75th	0.7322	0.4714	0.8289	0.5038	0.8091	0.4863	0.7382	0.7780
		90th	0.8553	0.2758	0.9199	0.2768	0.9194	0.2752	0.7152	0.7498
	Log-loss	25th	0.7148	0.8119	0.4368	0.8051	0.4468	0.8084	0.7170	0.8630
		Median	0.6959	0.6912	0.7020	0.7135	0.6784	0.7018	0.7558	0.6918
		Mean	0.6807	0.5648	0.7559	0.5978	0.7292	0.5803	0.7210	0.7503
		75th	0.7363	0.4788	0.8348	0.5251	0.8076	0.5003	0.7426	0.7679
		90th	0.8544	0.2551	0.9172	0.2476	0.9218	0.2500	0.7021	0.7522
100–149 kW	AUC	25th	0.7353	0.8112	0.4797	0.8408	0.4275	0.8253	0.6953	0.8736
		Median	0.6907	0.6878	0.6974	0.7093	0.6723	0.6968	0.7593	0.7041
		Mean	0.7148	0.6601	0.7518	0.6333	0.7721	0.6451	0.7757	0.7760
		75th	0.7713	0.5624	0.8492	0.5616	0.8434	0.5569	0.7944	0.7893
		90th	0.8973	0.5043	0.9446	0.4999	0.9414	0.4897	0.8513	0.8148
	Error	25th	0.7426	0.8181	0.4965	0.8422	0.4528	0.8295	0.7043	0.8663
		Median	0.6842	0.6855	0.6872	0.6908	0.6778	0.6861	0.7559	0.7173
		Mean	0.7128	0.6568	0.7519	0.6378	0.7658	0.6461	0.7833	0.7796
		75th	0.7630	0.5438	0.8403	0.5327	0.8422	0.5337	0.7946	0.7858
		90th	0.9000	0.5103	0.9418	0.4721	0.9476	0.4829	0.8472	0.8421
	Log-loss	25th	0.7342	0.8098	0.4729	0.8423	0.4203	0.8253	0.6993	0.8676
		Median	0.6926	0.6976	0.6926	0.6927	0.6925	0.6928	0.7632	0.7210
		Mean	0.7130	0.6583	0.7499	0.6310	0.7704	0.6431	0.7797	0.7712
		75th	0.7703	0.5533	0.8509	0.5652	0.8410	0.5537	0.7935	0.7972
		90th	0.9037	0.5338	0.9440	0.4907	0.9497	0.5028	0.8467	0.8371
150+ kW	AUC	25th	0.7209	0.8188	0.4557	0.8047	0.4773	0.8108	0.7776	0.8744
		Median	0.7053	0.7354	0.6850	0.6780	0.7342	0.7023	0.7661	0.7711
		Mean	0.6744	0.6223	0.7123	0.5740	0.7467	0.5916	0.7199	0.7690
		75th	0.7828	0.5725	0.8673	0.6210	0.8384	0.5906	0.8025	0.7651
		90th	0.8838	0.3653	0.9446	0.4583	0.9273	0.3965	0.7753	0.7950
	Error	25th	0.7403	0.8298	0.4990	0.8202	0.5077	0.8240	0.7905	0.8628
		Median	0.7033	0.7374	0.6813	0.6667	0.7422	0.6967	0.7724	0.7857
		Mean	0.6646	0.6082	0.7057	0.5696	0.7333	0.5834	0.7313	0.7423
		75th	0.7733	0.5520	0.8593	0.5983	0.8332	0.5713	0.7917	0.7848
		90th	0.8896	0.3976	0.9448	0.4583	0.9337	0.4097	0.7920	0.7961
	Log-loss	25th	0.7422	0.8351	0.4973	0.8175	0.5228	0.8250	0.7868	0.8761
		Median	0.6899	0.7255	0.6682	0.6514	0.7302	0.6819	0.7639	0.7939
		Mean	0.6784	0.6231	0.7145	0.5788	0.7500	0.5978	0.7209	0.7636
		75th	0.7868	0.5724	0.8752	0.6512	0.8333	0.6084	0.7994	0.7731
		90th	0.8896	0.3917	0.9450	0.4583	0.9337	0.4047	0.7693	0.8033

- touristic areas, where transient visitors create episodic but significant usage patterns.

Additional partial overlaps include:

- **financial**, shared between AUC and Log-loss,
- **supermarket**, shared between AUC and Error.

These confirm that practical, medium-duration errands (e.g., shopping) still play a meaningful role in driving demand for this power range.

3. Power range 150+ kW

For high-power charging stations, the following POIs are common across all three metrics (AUC, Error, Log-loss): **motorway_flag**, **amenity_facilities**, **house**, and **financial**.

These overlaps reflect the dual nature of HPC (High Power Charging) usage:

- **Transit-oriented behaviour:** the presence of **motorway_flag** confirms that HPC stations are strongly associated with long-distance travel corridors.
- **Local service-oriented behaviour:** features such as **amenity_facilities** and **financial** suggest that HPCs in urban or semi-urban areas benefit from integration with essential services.
- **Residential influence:** the recurrence of **house** indicates that HPC stations may also serve as substitutes for home charging in dense urban contexts.

A notable additional overlap includes:

- **leisure**, shared between AUC and Error,

suggesting that recreational areas can also constitute suitable locations for high-power charging, as charging may occur in parallel with leisure activities.

Synthesis Across All Ranges

A cross-range comparison reveals several high-level conclusions:

- Lower-power stations (50–99 kW) are associated with POIs tied to intermediate dwell times, such as fast-food, malls, and motorway proximity.
- Intermediate-power stations (100–149 kW) reflect more **local and community-driven demand**, with residential and touristic POIs emerging as key drivers.
- High-power stations (150+ kW) remain closely linked to **motorway infrastructure**, while also benefiting from nearby essential services and residential density.
- The partial overlaps between metrics (AUC, Error, Log-loss) indicate **robustness**: the observed POI importance patterns are not dependent on the evaluation metric.

Table 5.8: Percentage distribution of POIs by power range

	50-99	100-149	150+
house	40.41	36.98	26.39
shop	12.11	12.65	14.32
parking	10.81	10.72	20.72
amenity_facilities	5.13	5.81	6.27
leisure	5.09	4.88	3.09
sustenance_short	3.43	3.61	3.83
sustenance_long	3.15	3.63	4.45
sports	2.83	3.19	2.12
office	1.91	2.14	1.48
education	1.80	1.55	0.90
transportation	1.67	2.37	3.01
healthcare	1.55	1.55	1.25
tourism	1.34	1.88	0.70
public_service	1.24	1.08	0.83
financial	1.22	1.29	1.34
fuel	1.04	1.22	3.59
amenity_others	1.00	1.01	0.90
supermarket	0.99	0.99	0.97
spiritual	0.98	0.99	0.61
fast_food	0.95	0.99	1.48
entertainment	0.87	0.88	0.58
furniture	0.22	0.35	0.47
mall	0.19	0.17	0.50
department_store	0.05	0.05	0.17
wholesale	0.02	0.02	0.03

5.5.5 Discussion of POI Distributions Across Power Ranges

As shown in table 5.8, the percentage distributions of POIs associated with stations in the three power categories (50–99 kW, 100–149 kW, and 150+ kW) reveal several important patterns.

Similarity between 50–99 kW and 100–149 kW stations

A first notable result is the strong similarity in the POI distributions of the 50–99 kW and 100–149 kW stations. For most POI categories—including *house*, *shop*, *leisure*, *sustenance*, *sports*, *office*, *tourism*, and *public_service*—the percentages are very close across the two power classes.

This indicates that these two categories of chargers are typically deployed in similar spatial contexts, such as mixed-use or residential areas, and are therefore exposed to comparable local demand conditions. As a consequence, the differences observed in the predictive analysis between the two power ranges are unlikely to originate from the POI environment. Rather, they are more plausibly the result of differences in **operational usage**, such as dwell time, user profiles, or traffic patterns.

Distinct distribution for 150+ kW stations

In contrast, the POI distribution for 150+ kW stations diverges significantly from the other two ranges. Categories such as *parking* (increasing from about 10–11% to over 20%) and *fuel* (rising from around 1% to more than 3.5%) are markedly more frequent. Transportation-related POIs also gain importance.

This pattern reflects the functional role of high-power charging stations, which are primarily designed for short dwell-time charging. Their deployment typically favors accessibility-oriented locations such as large parking facilities, motorway corridors, fuel stations, and mobility hubs. This is consistent with modern HPC deployment practices and with user behaviour, as charging sessions at these stations typically last only 10–20 minutes.

Misalignment between POI frequency and POI importance

Interestingly, although POIs such as *parking* and *fuel* are highly represented in the 150+ kW distribution, they do not appear among the most important predictors of high usage in the feature-importance analysis. This suggests that these POIs describe where HPC sites are commonly built, but not what makes them highly used.

High utilization of HPC stations is instead more strongly driven by mobility-related factors, such as proximity to major road networks, accessibility, and the station's position within long-distance travel corridors.

Summary

Overall, the analysis reveals that:

- the 50–99 kW and 100–149 kW chargers share similar POI contexts, and differences in performance arise from different usage patterns rather than from different spatial environments;
- the 150+ kW chargers are deployed in structurally different contexts, consistent with fast-charging operational requirements;
- despite their high frequency, *parking* and *fuel* POIs do not explain high utilization, indicating that demand for HPCs is driven primarily by mobility flows rather than local amenities.

Chapter 6

Bilevel Design and Pricing of EV Charging Stations with Deviation-Flow

In this chapter, we propose a novel bilevel programming formulation that models the location, design, and pricing decisions for capacitated EV charging infrastructure. Our approach comprises several realistic features that enhance the applicability of the model to real-world urban settings. At the upper level, the leader, namely the Charging Point Operator (CPO), aims to maximize its overall revenue by deciding on the location, design, and pricing of stations, allowing for time-dependent pricing policies (e.g., peak vs. off-peak rates). The lower level describes the assignment of users to stations for minimizing the total charging cost through a deviation-flow based formulation, where users are allowed to make slight detours from their original path to access charging infrastructure and recharge at convenient costs. We assume that users have perfect knowledge of prices and availability at each station, reflecting the access to transparent real-time information systems provided by mobile applications. Moreover, the model accounts for location-specific attractiveness, representing user preferences as a percentage-based premium or discount on the willingness-to-pay price. To the best of our knowledge, there is no study in the literature that addresses the capacitated charging station location and pricing problem by jointly incorporating station design, multi-period pricing, and location-specific preferences.

The problem is formalized, and the bilevel formulation is presented in Section 1. Section 2 presents the computational setting and the experimental outcomes. Finally, Section 3 summarizes the key findings of the study and outlines perspectives.

6.1 Problem Definition and Bilevel Formulation

Consider a road network represented as a weighted directed graph $G = (N, A)$, where $N = \{1, \ldots, i\}$ is the set of nodes and A is the set of arcs. A subset $I \subseteq N$ of candidate nodes is given as the potential location of charging stations, and a discrete time horizon T is defined. Multiple connector types, distinguished by their power output, are included in the set K. EV drivers travel between origin-destination (OD) pairs within this network.

Given a set of EV flow demands Q, each $q \in Q$ gathers a set of customers with homogeneous economic and qualitative preferences, formally described by the parameter tuple $\langle o_q, d_q, r_q, \delta_q, k_q, f_q, t_q, \theta_q, h_q \rangle$:

- o_q, d_q denote the origin and destination nodes in I, respectively;
- r_q indicates the range of vehicles q on a full charge;
- δ_q represents the maximum deviation distance that drivers q are willing to tolerate;
- k_q is the connector type in K required by demand q;
- f_q gives the volume of flow demand q;
- t_q represents the origin/service period in T of demand q;
- θ_q is the unitary price users q are willing to pay;
- h_q energy requirement (in kWh) by demand q.

Both the OD paths, in the absence of charging needs, and the deviation paths, used to reach the destination going via a charging station, are built by computing the relevant shortest paths. The setting allows for multiple vehicle flows, that is, different demands, between the same OD pair, while accounting for different detour tolerances. Furthermore, we assumed that an EV departs from its origin with a half-full battery and must arrive at its destination with at least a half-full battery, as proposed in [26, 57]. To complete the mathematical notation, Table 6.1 is presented.

6.1.1 EV Users' Level

The lower level $\mathcal{L}$ describes the decision-making process of EV users choosing a charging station. Let $\bar{\mathbf{z}}$ and $\bar{\mathbf{p}}$ be respectively the design and pricing decisions taken by the CPO. The formulation reads as:

$$[\mathcal{L}] \qquad \min_{y} \sum_{q \in Q} \sum_{i \in I_q} \left[h_{qi} \bar{p}_{ik_q}^{t_q} - h_q \theta_q (1 + a_{qi}) \right] y_{qi} \tag{6.1}$$

$$\text{s.t.} \sum_{q \in Q_{ik}^t} y_{qi} \le \bar{z}_{ik}, \qquad \forall i \in I, k \in K, t \in T \tag{6.2}$$

$$\sum_{i \in I_q} y_{qi} \le f_q, \qquad \forall q \in Q \tag{6.3}$$

$$y_{qi} \ge 0 \qquad \forall q \in Q, i \in I_q \tag{6.4}$$

The objective function (6.1) minimizes the total difference between the actual charging costs, $h_{qi} \bar{p}_{ik_q}^{t_q}$, that is, the price paid by users served y_{qi} at station i with connector type

Table 6.1: Notation summary

Sets	**Description**
I_q	Set of candidate locations that can serve demand q.
Q_{ik}^t	Demands $q \in Q$ requiring connector k at facility $i \in I_q$ in period t.
Parameters	**Description**
B	Available budget for investments.
b_i	Cost to build a station at location i.
c_k	Unitary cost to install a connector of type k.
U_i	Maximum number of connectors allowed at station i.
a_{qi}	Attractiveness of station i according to demand q.
h_{qi}	Recharge amount in kWh at station i demanded by q.
σ_{qi}	Number of periods to charge a type q demand at station i.
Variables	**Description**
x_i	1 if location i is selected, 0 otherwise.
z_{ik}	Number of connectors of type k installed at location i.
p_{ik}^t	Price per kWh for connector k at station i during period t.
y_{qi}	Part of flow demand q served by location i.

k_q during period t_q, and the users' maximum willingness to pay, $h_q\theta_q$, adjusted by the station attractiveness factor a_{qi}. If the former cost exceeds the latter at station i, then minimization forces $y_{qi} = 0$, meaning demand q will not use station i. Conversely, if the willingness-to-pay is respected, the flow demand q may be captured. This may occur only in its origin period t_q due to the definition of y_{qi}. The amount h_{qi} incorporates the base supply need h_q and the charge consumption for deviating to location i.

Inequalities (6.2) ensure that the total demand assigned to station i does not exceed the available number of connectors of type k in each period t. In particular, the subset $Q_{ik}^t = \{q \in Q |\ i \in I_q, k_q = k, t_q \le t \le t_q + \sigma_{qi} - 1\}$ comprises flow demands q reaching location i and occurring in t_q such that period t is needed for recharging. In particular, σ_{qi} is the corresponding number of charging periods required by user vehicles q before releasing the connectors. Inequalities (6.3) impose that the total number of vehicles for each q assigned to different charging stations cannot exceed the flow volume f_q. Finally, (6.4) defines the flow demand variables y_{qi} as continuous. In preliminary tests on over 3000 instances, enforcing integrality on y_{qi} affected the optimal solution in roughly 0.5% of cases, with an integrality gap below 0.05%.

6.1.2 CPO's Level

The upper level $\mathcal{U}$ models the investments and pricing decisions of the CPO:

$$[\mathcal{U}] \qquad \max_{x,z,p,y} \sum_{q \in Q} \sum_{i \in I_q} h_{qi} p_{ik_q}^{t_q} y_{qi} \tag{6.5}$$

$$\text{s.t.} \sum_{i \in I} \left(b_i \cdot x_i + \sum_{k \in K} c_k \cdot z_{ik} \right) \leq B \tag{6.6}$$

$$\sum_{k \in K} z_{ik} \leq U_i x_i \qquad \forall i \in I \tag{6.7}$$

$$x_i \in \{0, 1\} \qquad \forall i \in I \tag{6.8}$$

$$z_{ik} \geq 0, \text{ integer} \qquad \forall i \in I, k \in K \tag{6.9}$$

$$p_{ik}^t \geq 0 \qquad \forall i \in I, k \in K, t \in T \tag{6.10}$$

$$\mathbf{y} \in \mathcal{L} \tag{6.11}$$

The bilinear objective function (6.5) sums up the revenues collected from the flow demand served y_{qi}, each for a power amount h_{qi}, at price $p_{ik_q}^{t_q}$. Inequality (6.6) ensures that the total investment cost, given by location decisions x_i and design choices z_{ik}, does not exceed the available budget B. Constraint (6.7) ensures that the total number of installed connectors does not exceed the structural limit U_i at each selected location. Finally, (6.8)-(6.10) specify the variables' domain, and (6.11) sets flow demand variables to be the optimal solution of the lower level.

The pricing scheme can incorporate wholesale electricity costs (including a sustainability markup) through modifications to either constraints (6.10) or objective (6.5). These costs typically depend on the infrastructure configuration, power agreements, and grid-decoupling technologies like battery storage.

6.1.3 Bilinear Single-level Optimisation Model

Given the boundedness and non-emptiness of the lower level $\mathcal{L}$, the bilevel model can be restated as a single-level formulation by using the primal-dual reformulation discussed in [33], obtaining the following equivalent single-level problem:

$$\max_{x,z,p,y} \sum_{n \in N} \sum_{q \in Q^n} \sum_{i \in I_q} h_{qi} p_{in}^{t_q} y_{qi} \tag{6.12}$$

$$\text{s.t.} \sum_{i \in I} \left(b_i \cdot x_i + \sum_{n \in N} c_n \cdot z_{in} \right) \leq B, \tag{6.13}$$

$$\sum_{n \in N} z_{in} \leq N_i x_i, \qquad \forall i \in I, \tag{6.14}$$

$$p_{in}^t \geq 0, \qquad \forall i \in I, \forall n \in N, \forall t \in T, \tag{6.15}$$

$$z_{in} \in \mathbb{N}, \qquad \forall i \in I, \forall n \in N, \tag{6.16}$$

$$h_{qi}\bar{p}^{t_q}_{in_q} \leq h_q\theta_q(1 + a_{qi}) + P^{t_q}_{in_q} w_{qi}, \qquad \forall q \in Q, i \in I_q \tag{6.17}$$
$$y_{qi} \leq f_q(1 - w_{qi}), \qquad \forall q \in Q, i \in I_q \tag{6.18}$$
$$w_{qi} \in \{0, 1\}, \qquad \forall q \in Q, i \in I_q \tag{6.19}$$

$$\sum_{q \in Q^t_{i,n}} y_{qi} \leq z_{in}, \qquad \forall i \in I, \forall n \in N, \forall t \in T, \tag{6.20}$$
$$\sum_{i \in I_q} y_{qi} \leq f_q, \qquad \forall q \in Q, \tag{6.21}$$
$$y_{qi} \geq 0, \qquad \forall i \in I, \forall q \in Q. \tag{6.22}$$

$$\sum_{t=t_q}^{t_q+\sigma_{qi}-1} \alpha^t_{in} + \beta_q \leq h_{qi}p^{t_q}_{in_q} - h_q\theta_q(1 + a_{qi}), \qquad \forall q \in Q, \forall i \in I_q, \tag{6.23}$$
$$\alpha^t_{in} \leq 0, \qquad \forall i \in I, \forall n \in N, \forall t \in T, \tag{6.24}$$
$$\beta_q \leq 0, \qquad \forall q \in Q. \tag{6.25}$$

$$\alpha^t_{in} \cdot \left(z_{in} - \sum_{q \in Q^t_{i,n}} y_{qi} \right) = 0, \qquad \forall i \in I, \forall n \in N, \forall t \in T, \tag{6.26}$$
$$\beta_q \cdot \left(f_q - \sum_{i \in I_q} y_{qi} \right) = 0, \qquad \forall q \in Q. \tag{6.27}$$

Constraints (6.17), (6.18), and (6.19) derive from the second-level objective function, which implicitly assumes that, for any $q \in Q$ and $i \in I_q$:

$$\text{if } h_{qi}\bar{p}^{t_q}_{in_q} > h_q\theta_q(1 + a_{qi}), \text{ then } y_{qi} \leq 0. \tag{6.28}$$

This condition ensures that users will only choose a charging station if the cost does not exceed their willingness to pay. Since y_{qi} must remain non-negative, this requirement is enforced explicitly through the aforementioned constraints.

The upper bound $P^{t_q}_{in_q}$ is defined as:

$$P^{t_q}_{in_q} = \max_{l \in Q_i : t_l = t_q, n_l = n_q} \{h_{li}\theta_l(1 + a_{li})\} - h_q\theta_q(1 + a_{qi}). \tag{6.29}$$

This formulation guarantees that the price for a full charge h_{qi} at station i during period t never exceeds the highest willingness to pay for a full charge among all demands l using station i at the same time t. If this condition were violated, the station i would be automatically excluded from the market during that period, as no user would choose it.

Constraints (6.23), (6.24), and (6.25) are the constraints associated with the dual problem of the second-level problem.

Constraints (6.26) and (6.27) are the complementary slackness constraints.

6.2 Computational Experiments

Numerical experiments were done to assess the performance of the single-level formulation on the 25-node road benchmark network in Fig. 6.1, see [47]. The formulation was implemented in Python and solved by using Gurobi Optimizer 12.0.0. All experiments were performed on a machine equipped with AMD Ryzen 7 4800H and 16 GB of RAM.

Figure 6.1: The 25-node road network. The labels on the edges represent travel distances.

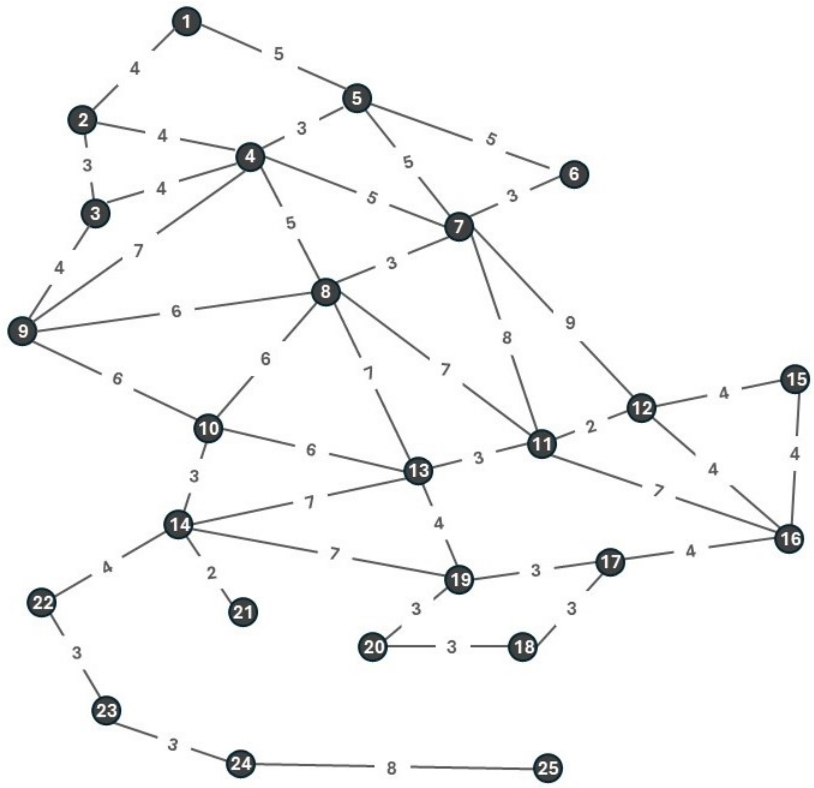

6.2.1 Instance Generation

25-node graph

We generated a collection of synthetic instances simulating EV travel demand over the network described above. First, a pool of 125 OD pairs was sampled from the set of network nodes. Three levels of demand were defined such that $|Q| \in \{500, 750, 1000\}$. For each demand level, we created three distinct instances by independently sampling $|Q|$ OD pairs from the pool, allowing the same OD pair to be selected multiple times within each instance (i.e., sampling with replacement). All instances simulate a total flow of vehicles of 4000 units, randomly distributed among the $|Q|$ demands.

Each $q \in Q$ was assigned a connector power level $k_q \in \{22, 50, 150\}$ kW, randomly sampled from a uniform distribution over the available connector types. Then, for each length of time horizon $|T| \in \{16, 32, 48\}$, where each period represents a 15-minute interval, time origins t_q were randomly extracted in $t_q \in \{0, 1, \ldots, T-1\}$. The vehicle ranges and deviations were treated as constant scenario parameters shared among all the demands of an instance. In particular, we generated instances for all combinations of $r_q \in \{6, 12, 18\}$ and $\delta_q \in \{0, 9, 12\}$.

Finally, the number of candidate charging station locations was randomly extracted, with $|I| \in \{2, 5, 10, 15, 20\}$. Altogether, 1215 instances were obtained.

Parameters θ_q, h_q and σ_{qi} were estimated from prices, connector features and power demand volumes derived from real charging session data, comprising per-session duration and connector type, available at the partner company's database [5]. The attractiveness factors a_{qi} were randomly sampled from the uniform interval [–0.02, 0.08]. Lastly, the budget was fixed to B = 500,000 and b_i were randomly sampled from [30,000, 50,000]. Station capacities U_i were drawn uniformly as an integer between 2 and 10, whereas c_k were set to 200, 300, and 400, respectively for 22 kW, 50 kW, and 150 kW connectors.

Realistic data

This section investigates the computational behaviour of the proposed optimisation model on the set of realistic instances generated on the subnetwork of the city of Milan. The underlying transportation network consists of 1,310 nodes and 1,608 directed arcs and represents a geographically coherent and operationally relevant portion of the urban road infrastructure.

The generation process of the total demand flows is described in detail in Section 3.6. Then, each dataset is constructed by randomly assigning a range (30%, 50%, 60% of the shortest path) and a deviation parameter (5%, 10%, 15% of the shortest path) to each origin–destination flow, and retaining only those flows whose shortest path length exceeds 3 km. This filtering procedure results in a pool of 9,608 feasible flows. From this pool, three independent datasets are created by randomly sampling 6,000, 7,000, and 8,000 flows, respectively. For each dataset, three independent instances are generated in order to account for variability induced by the random sampling process.

Regarding the modelling parameters, the attractiveness coefficients a_{qi} are derived from the machine learning analysis presented in Chapter 5, which quantifies the influence of spatial and contextual features on users' charging choices. The parameters h_q, representing the demand-related attributes of the flows, are obtained from data provided by Atlante and are assigned randomly across flows to preserve realism while avoiding structural bias.

The experimental setup considers different problem sizes by varying the number of candidate charging locations and the number of observations. Specifically, the number of observations is set to $\{6000, 7000, 8000\}$, while the number of candidate locations is selected from the set $\{20, 50, 100, 150, 200\}$. A total of three instances is generated for each configuration. The budget is fixed to $B = 500{,}000$, and the installation costs c_n are defined as a function of the number of candidate locations, with values increasing proportionally to reflect higher infrastructure complexity.

This experimental design allows us to assess the scalability and robustness of the model under conditions that closely resemble real-world urban planning scenarios, while maintaining sufficient variability to stress-test the optimisation framework.

6.2.2 Computational Results

25-node graph

3600 seconds time limit. Table 6.2 shows the model performance on each group $G_{|I|,|T|}$ and $|Q|$ pair. Each entry refers to 27 instances, that is 3 instances per r_q and δ_q combination, sharing the same number of candidate locations and periods. In particular, the average CPU time (t_{avg}) and the maximum CPU time (t_{max}) are reported. The average optimality gap (in %) and the number of timed-out instances (in brackets) are reported in bold within the latter column when the time limit occurred. Finally, the overall average and total CPU times are given. Notice that the minimum CPU time has been omitted, being at most 0.26 seconds across all instances.

The results in Table 6.2 show that the solution time grows proportionally with the number of candidate locations and total demand. While instances with low values of $|I|$ (e.g., $G_{2,16}$ or $G_{2,48}$) are solved very quickly even for a total demand of 1000, the CPU time significantly grows with higher demand amounts and for a wider number of candidate locations. Starting from $G_{15,16}$ with $|Q| = 750$, the time limit occurred in 34 instances, corresponding to 2.80% of the cases. Nevertheless, optimality gaps remained small, at most 0.22% in group $G_{20,32}$ and $|Q| = 1000$.

Although the solution time generally increases with the cardinality of I, interestingly, its growth is not strictly linear with respect to either $|I|$ or $|T|$, suggesting that other factors contribute to the practical hardness. For example, group $G_{15,32}$ ($G_{20,16}$) with $|Q| = 500$ exhibits particularly high solution times despite not having the highest number of candidate locations (periods). This behavior is likely related to congestion effects: as the number of time periods increases, demands are more evenly distributed over time, which helps mitigate demand peaks and, consequently, reduces the practical difficulty. On the other hand, a larger number of time periods required to complete a charging session (σ_{qi}) can further amplify congestion for smaller $|T|$, increasing competition for access to infrastructure and making the instances challenging. Indeed, when comparing instances with

Table 6.2: Average and maximum CPU times (in seconds) per instance group $G_{|I|,|T|}$. When the time limit is reached, the percentage average optimality gap and the number of unsolved instances (in brackets) are reported in columns t_{max}.

Group	$\|Q\| = 500$		$\|Q\| = 750$		$\|Q\| = 1000$	
	t_{avg}	t_{max}	t_{avg}	t_{max}	t_{avg}	t_{max}
$G_{2,16}$	0.23	1.06	0.33	1.78	0.40	3.65
$G_{2,32}$	0.22	1.67	0.34	1.94	0.73	11.75
$G_{2,48}$	0.20	1.20	0.50	3.80	0.59	7.54
$G_{5,16}$	1.09	9.88	2.78	15.57	6.07	63.09
$G_{5,32}$	1.17	14.25	4.34	63.43	10.69	86.91
$G_{5,48}$	0.54	2.14	3.26	55.88	13.20	180.59
$G_{10,16}$	3.65	43.69	39.14	336.29	24.45	162.86
$G_{10,32}$	4.94	51.91	31.71	239.13	55.72	454.28
$G_{10,48}$	3.13	38.28	41.32	553.88	35.90	246.88
$G_{15,16}$	25.18	154.84	483.40	**0.05 (3)**	615.46	**0.04 (3)**
$G_{15,32}$	82.28	1638.50	377.16	**0.01 (1)**	600.87	**0.02 (3)**
$G_{15,48}$	17.83	146.27	185.46	1658.47	300.15	**0.03 (1)**
$G_{20,16}$	132.19	1839.95	814.52	**0.09 (5)**	809.04	**0.09 (5)**
$G_{20,32}$	79.90	785.84	599.50	**0.03 (3)**	684.38	**0.09 (4)**
$G_{20,48}$	98.73	933.54	313.20	**0.03 (1)**	803.23	**0.07 (5)**
Avg.	**30.09**	-	**193.13**	-	**264.06**	-
Total	**12184.70**	-	**78217.95**	-	**106943.64**	-

the same number of stations, available locations, and demand level, those with $t = 48$ often result in lower solution times than those with $t = 16$ or $t = 32$. Moreover, the computational time demanded by the three instances with the same parameter configuration may vary significantly, highlighting how the distribution of demands in time can strongly affect the performance. For example, the CPU time asked by the group $G_{20,48}$, $|Q| = 1000$ with $r_q = 18$ and $\delta_q = 9$ range between 25.73 and 150.72 seconds.

Figure 6.2 shows the relationship between solution time (in logarithmic scale) and the combinations of vehicle range r_q and deviation tolerance δ_q, for different numbers of time intervals $|T|$. As expected, we observe that increasing the range r_q leads to a general increase in solution time due to longer-distance demands. The impact of detour tolerance δ_q appears less pronounced. More specifically, instances with $\delta_q = 9$ and $\delta_q = 12$ exhibit similar computational behavior across all ranges, suggesting that the problem is not sensitive to limited variations in the deviation.

300 seconds time limit. To assess the computational performance of the proposed optimisation model, we conducted a set of experiments in which a strict time limit was imposed on the solver. In particular, a timeout of 5 minutes was enforced for each instance in order to analyse the behaviour of the optimality gap under limited computational resources. This choice reflects a realistic operational setting, where solution times are necessarily bounded, and allows us to evaluate the trade-off between solution quality and computational effort. For each instance reaching the time limit, the final relative optimality gap

Figure 6.2: Log-transformed CPU times for each instance, grouped by (r_q, δ_q) combinations.

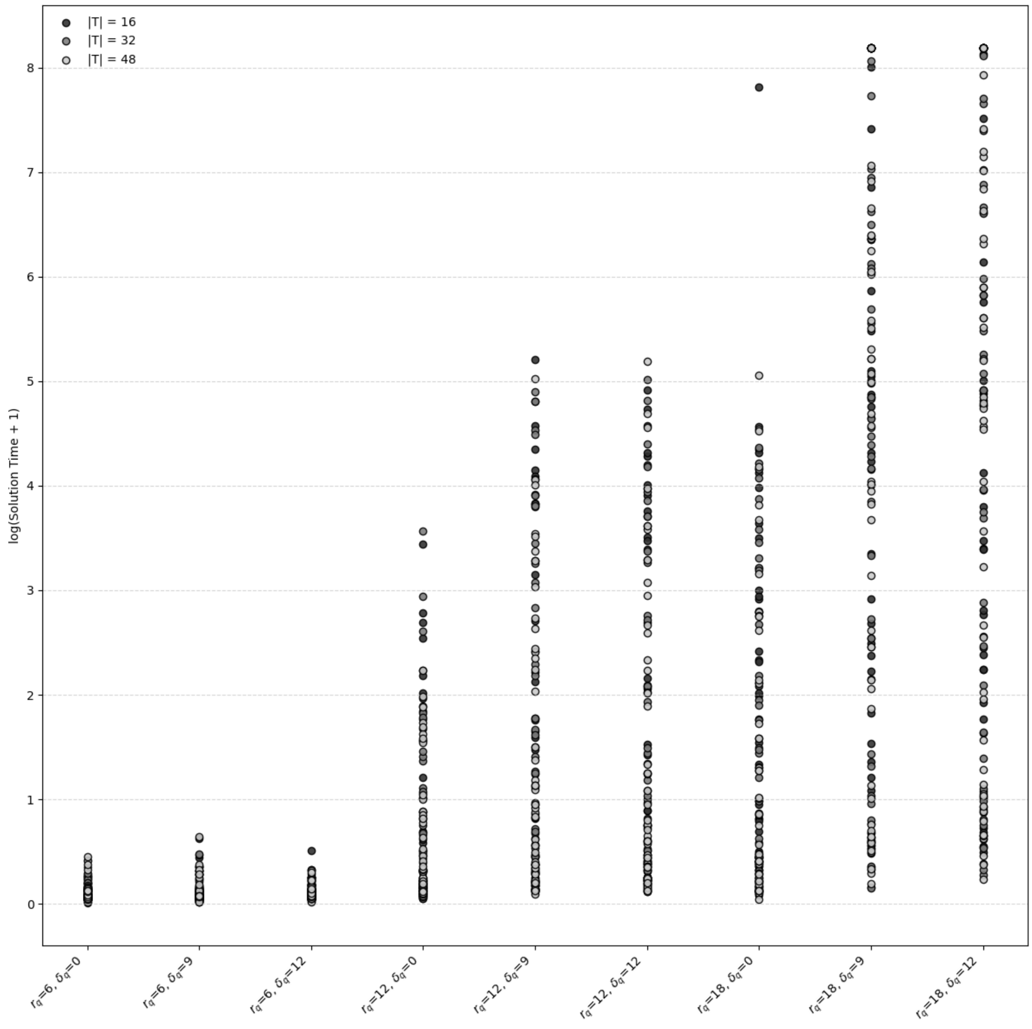

reported by the solver was recorded and used as a performance indicator. The analysis of these gaps provides insights into the scalability of the model and highlights the instance characteristics that most strongly affect convergence within a fixed time horizon.

Table 6.3: Instances with 3600-second solver timeout.

Num_obs	Time	Range	Dev	Istance	Stations	Gap
750	16	18	9	1	20	0.811
750	16	18	9	2	15	1.684
750	16	18	9	2	20	1.710
750	16	18	9	3	20	0.432
750	16	18	12	1	15	0.138
750	16	18	12	1	20	0.932
750	16	18	12	2	20	0.885
750	16	18	12	3	15	0.490
750	32	18	9	2	20	1.717
750	32	18	12	1	15	0.331
750	32	18	12	1	20	0.748
750	32	18	12	3	20	1.359
750	48	18	12	3	20	1.155
1000	16	18	9	1	20	1.613
1000	16	18	9	2	20	0.554
1000	16	18	9	3	15	0.221
1000	16	18	12	1	15	0.826
1000	16	18	12	1	20	0.268
1000	16	18	12	2	20	0.491
1000	16	18	12	3	15	0.212
1000	16	18	12	3	20	2.003
1000	32	18	9	1	20	2.646
1000	32	18	9	2	15	0.667
1000	32	18	9	2	20	0.623
1000	32	18	12	1	20	1.392
1000	32	18	12	2	15	0.400
1000	32	18	12	3	15	0.105
1000	32	18	12	3	20	0.894
1000	48	18	9	1	20	0.394
1000	48	18	9	2	20	0.330
1000	48	18	9	3	20	0.001
1000	48	18	12	1	15	0.223
1000	48	18	12	1	20	1.306
1000	48	18	12	3	20	0.856

Table 6.3 reports only the instances that reached the imposed time limit when a timeout of 1 hour was applied. These instances represent the most computationally challenging configurations, for which the solver was unable to prove optimality within the allotted time. By focusing exclusively on this subset, the table highlights the parameter combinations that lead to increased solution difficulty and persistent optimality gaps, providing a clearer understanding of the conditions under which the model exhibits scalability limitations.

Despite the strict time limit of only 5 minutes, the solver was able to achieve very small optimality gaps for all of the tested instances. This result indicates that the proposed formulation exhibits a strong convergence behaviour, allowing high-quality solutions to be identified early in the search process. In many cases, near-optimal solutions are obtained well before the time limit is reached, suggesting that additional computational time primarily serves to certify optimality rather than to significantly improve the objective value. These findings are particularly relevant from a practical perspective, as they demonstrate that the model can deliver reliable solutions within short time horizons, even for instances that are challenging when longer time limits are considered.

Realistic data results

Table 6.4: Average and maximum CPU times (in seconds) per instance group $G_{|I|}$. Only instances solved within the time limit are included.

Group	$\lvert Q\rvert = 6000$		$\lvert Q\rvert = 7000$		$\lvert Q\rvert = 8000$	
	t_{avg}	t_{max}	t_{avg}	t_{max}	t_{avg}	t_{max}
G_{20}	53.37	181.42	14.71	44.72	25.33	57.73
G_{50}	176.02	920.06	381.18	798.69	524.11	1369.54
G_{100}	658.04	1565.52	961.32	1489.21	901.37	1661.62
G_{150}	1262.03	1653.70	505.09	505.09	1568.34	1568.34
Avg.	**351.89**	-	**418.12**	-	**472.75**	-
Total	**9852.87**	-	**10034.97**	-	**10873.28**	-

Table 6.4 reports the average and maximum CPU times computed over the instances that were solved to optimality within the imposed time limit of 1800 seconds. The results highlight a clear and systematic increase in computational effort as the number of candidate locations grows, confirming the expected impact of problem size on solution times.

The instance groups reported in Table 6.4 are defined according to the number of candidate charging locations considered in the optimisation problem. Specifically, each group $G_{|I|}$ collects all instances sharing the same cardinality $|I|$ of the set of potential locations, independently of the number of demand flows or the dataset realization.

This grouping strategy allows us to isolate the impact of the size of the location decision space on the computational performance of the model. By aggregating instances with identical values of $|I|$, the reported statistics capture how the number of candidate sites alone influences solution times, while averaging out the variability induced by different demand realizations and dataset samples.

For each group and each value of $|Q|$, the average and maximum CPU times are computed over all instances that were solved to optimality within the prescribed time limit. Instances reaching the time limit without proving optimality are excluded from this analysis and are examined separately through the study of optimality gaps.

For small and medium-sized groups, such as G_{20} and G_{50}, the model exhibits very good computational performance, with average solution times remaining well below a few minutes across all values of $|Q|$. In these cases, even the maximum observed times are moderate, indicating a stable and reliable convergence behaviour.

As the number of candidate locations increases to G_{100} and G_{150}, the computational burden becomes significantly higher. Average solution times exceed several hundred seconds and, in some configurations, approach the time limit. This effect is particularly evident for larger demand sets, where the interaction between the number of flows and the expanded decision space leads to a marked increase in complexity.

Table 6.5 reports the average and maximum optimality gaps for the instance groups, considering only runs for which the final gap is below 50%. The results show that, for small and medium-sized groups (G_{20}, G_{50}, and G_{100}), the solver consistently achieves optimal

Table 6.5: Average and maximum optimality gaps (in %) per instance group $G_{|I|}$. Only runs with an optimality gap below 50% are included.

Group	$\lvert Q\rvert = 6000$		$\lvert Q\rvert = 7000$		$\lvert Q\rvert = 8000$	
	gap_{avg}	gap_{max}	gap_{avg}	gap_{max}	gap_{avg}	gap_{max}
G_{20}	0.00	0.00	0.00	0.00	0.00	0.00
G_{50}	0.00	0.00	0.01	0.12	0.02	0.15
G_{100}	0.00	0.02	0.00	0.01	0.00	0.01
G_{150}	0.79	2.19	0.99	3.46	0.60	1.82
G_{200}	0.96	1.62	0.70	0.92	0.73	1.26
Avg.	**0.31**	-	**0.21**	-	**0.20**	-
Total	**13.09**	-	**7.47**	-	**7.41**	-

or near-optimal solutions, with both average and maximum gaps remaining close to zero across all values of $|Q|$. This indicates a strong convergence behaviour of the proposed formulation when the size of the location decision set is limited.

As the number of candidate locations increases, a moderate deterioration in solution quality can be observed. For groups G_{150} and G_{200}, average gaps remain below 1%, while maximum gaps stay within a few percentage points. Although these instances are computationally more challenging, the solver is still able to identify high-quality solutions within the prescribed time limit.

Overall, the gap analysis suggests that most of the computational difficulty associated with larger instances is related to the certification of optimality rather than to the identification of good feasible solutions. From a practical standpoint, the small gaps observed even for large-scale configurations support the applicability of the proposed model in realistic urban planning contexts, where near-optimal solutions obtained within limited computational time are often sufficient for decision-making.

Table 6.6: Instances with large optimality gaps at the time limit.

Num_obs	Temp	Dataset	Inst.	Stations	Gap (%)
6000	80	1	1	200	249.43
6000	83	1	3	150	446.77
6000	83	3	2	200	427.56
7000	83	1	1	200	148.78
7000	83	1	3	200	552.58
7000	83	2	1	200	253.62
7000	83	2	2	200	∞
7000	83	2	3	150	366.51
7000	83	2	3	200	51.44
7000	83	3	1	200	787.07
7000	83	3	2	150	482.40
7000	83	3	3	150	151.04
7000	83	3	3	200	610.25
8000	80	1	1	150	285.72
8000	80	1	1	200	788.86
8000	83	1	2	200	388.72
8000	83	2	1	150	865.53
8000	83	2	2	200	442.66
8000	83	2	3	150	132.01
8000	83	3	1	200	∞

Table 6.6 reports the instances that reached the imposed time limit and for which the

solver was unable to significantly reduce the optimality gap. These cases correspond to the most computationally challenging configurations and are predominantly associated with large numbers of candidate locations and high demand levels.

The reported gaps are often very large, and in some instances unbounded, indicating that the solver struggles not only to certify optimality but also to substantially improve the incumbent solution within the allotted time. This behaviour suggests that, for these instances, the size of the feasible region and the combinatorial nature of the location decisions severely hinder the exploration of the solution space.

It is worth noting that such extreme gaps arise only in a limited subset of configurations and typically occur when both the number of observations and the number of candidate locations are large.

6.3 Conclusions and Perspectives

In this chapter, we consider the combined location and pricing problem of electric vehicle (EV) charging stations. We propose a novel bilevel mixed-integer bilinear programming formulation, where the upper level models the charging point operator's (CPO's) decision-making process and the lower level captures users' supply choices. Specifically, the upper level selects the locations and designs of stations, together with the service prices for each connector type and time period, in order to maximize the operator's profit under a budget constraint. The lower level describes the assignment of flow demands to stations that minimizes users' total recharging costs, accounting for users' willingness to pay and the stations' attractiveness.

Computational experiments were conducted on both randomly generated instances and on more realistic instances designed to reflect key characteristics of real-world charging networks. Random instances were first solved with a one-hour time limit. Since the resulting optimality gaps were consistently very small, we subsequently reduced the time limit to five minutes to assess whether comparable solution quality could still be achieved under a tighter computational budget. Realistic instances, which incorporate empirically motivated demand and network features, were instead solved using a time limit of thirty minutes.

The computational analysis highlights that network congestion—largely driven by temporal demand distributions and charging durations—plays a critical role in determining the practical hardness of the instances. In particular, highly congested scenarios significantly increase solution times and often prevent the solver from closing the optimality gap within the imposed time limits, especially for realistic instances.

As future perspectives, we aim to extend the computational campaign to larger and more detailed real-world urban networks, incorporating richer demand patterns, traffic dynamics, and pricing schemes. Moreover, we plan to enhance the solution methodology by investigating advanced algorithmic strategies, such as dynamic constraint handling techniques (e.g., Benders' decomposition and branch-and-cut methods), as well as alternative bilevel reformulation approaches that may further improve scalability and computational performance.

Chapter 7

Conclusions and perspectives

The rapid diffusion of electric vehicles (EVs) represents a key opportunity to reduce emissions and improve the sustainability of urban mobility systems. However, the success of this transition is tightly linked to the availability of a charging infrastructure that is not only widespread but also economically viable and aligned with user behavior. This thesis addressed this challenge by developing an integrated analytical framework that combines data-driven demand analysis with an advanced optimization model to support strategic decision-making in EV charging infrastructure planning.

The first part of this work focused on understanding and predicting charging demand using real-world data. By constructing a rich dataset that integrates spatial, traffic, and infrastructural features, a machine learning–based classification model was developed to identify locations likely to experience high charging activity. The results demonstrate that the proposed model is able to predict high-demand locations with an accuracy of approximately 70%, offering a practical screening tool for candidate sites. Beyond predictive performance, the analysis provided valuable insights into the determinants of charging station usage, highlighting the importance of connector power, traffic intensity, and surrounding points of interest. Importantly, user preferences were inferred directly from observed charging session data rather than relying solely on stated preferences from surveys, ensuring a behaviorally grounded representation of demand.

Building upon these empirical findings, the second and main contribution of the thesis is the formulation of a novel bilevel optimization model for the capacitated location and pricing of EV charging stations. The upper-level problem represents the decisions of a charging point operator, including station location, station design in terms of the number of connectors, and the selection of time-dependent pricing policies. The lower-level problem captures user behavior through a deviation flow capturing approach, in which drivers choose charging stations by minimizing individual charging costs while accounting for prices, travel distance, station availability, and location-specific attractiveness. Notably, the attractiveness component of the utility function is informed by the machine learning analysis, thereby directly linking observed charging behavior to the optimization framework.

Computational experiments conducted on both randomly generated and realistic instances demonstrate the viability of the proposed approach and provide insights into its computational behavior. Random instances were initially solved with a one-hour time limit, yielding very small optimality gaps; the time limit was then reduced to five minutes to assess the robustness of solution quality under tighter computational constraints. Realistic instances, constructed to reflect key features of real-world charging networks, were instead solved using a time limit of thirty minutes. The results highlight the significant impact of temporal demand distributions and charging durations on computational complexity and solution structure. In particular, capacity constraints and time-dependent pricing play a crucial role in shaping user flows and station profitability, while neglecting demand response may lead to suboptimal or even infeasible infrastructure designs. To the best of our knowledge, this work represents the first optimization model that jointly addresses capacitated charging station location, station design, multi-period pricing, and spatially heterogeneous user preferences inferred from real charging data.

Overall, this thesis contributes both methodologically and practically to the planning of EV charging infrastructure. By integrating machine learning and bilevel optimization within a unified framework, it provides decision-makers with tools to better anticipate demand, evaluate investment strategies, and design charging networks that are efficient, profitable, and user-oriented. Future research may extend this framework by incorporating uncertainty in demand and prices, renewable energy integration, energy storage systems, and competitive interactions among multiple charging service providers.

Bibliography

[1] M. Adil, M. P. Mahmud, A. Kouzani, and S. Y. Khoo. A stackelberg game approach to find optimal location and pricing for network of charging stations. In *2023 IEEE IAS Global Conference on Emerging Technologies (GlobConET)*, pages 1–6, 2023.

[2] T. Akiba, S. Sano, T. Yanase, T. Ohta, and M. Koyama. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2623–2631. ACM, 2019.

[3] Alternative Fuels Observatory. European alternative fuels observatory – consumer monitor. `https://alternative-fuels-observatory.ec.europa.eu/consumer-portal/consumer-monitor`, 2023. Accessed April 14, 2025.

[4] O. Arslan and O. E. Karaşan. A benders decomposition approach for the charging station location problem with plug-in hybrid electric vehicles. *Transportation Research Part B: Methodological*, 93:670–695, 2016.

[5] Atlante. Atlante – fast and ultra-fast charging network for southern europe. `https://atlante.energy/it/`, 2025. Accessed June 5, 2025.

[6] Automobile Club d'Italia (ACI). Aci open data – registered vehicles by region. `https://aci.gov.it/attivita-e-progetti/studi-e-ricerche/open-data/`, 2023. Accessed June 5, 2025.

[7] I. Baffo, S. Minucci, S. Ubertini, G. Calabrò, A. Genovese, and N. Andrenacci. Optimal sizing and location of a fast charging infrastructure network for urban areas. In *2020 AEIT International Annual Conference (AEIT)*, pages 1–6, 2020.

[8] J. F. Bard. *Practical Bilevel Optimization: Algorithms and Applications*. Kluwer Academic Publishers, Dordrecht, 1998.

[9] A. Ben-Tal and A. Nemirovski. Duality and convex programming. *SIAM Review*, 42(4):469–518, 2000.

[10] L. Breiman. Random forests. *Machine Learning*, 45(1):5–32, 2001.

[11] P. Calamai and L. N. Vicente. A review of bilevel programming: A taxonomy and a bibliography. *European Journal of Operational Research*, 69(2):190–211, 1994.

[12] T. Chen and C. Guestrin. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 785–794. ACM, 2016.

[13] S. H. Chung and C. Kwon. Multi-period planning for electric car charging station locations: A case of korean expressways. *European Journal of Operational Research*, 242(2):677–687, 2015.

[14] B. Colson, P. Marcotte, and G. Savard. An overview of bilevel optimization. *Annals of Operations Research*, 153(1):235–256, 2007.

[15] M. Cruz-Zambrano, C. Corchero, L. Igualada-Gonzalez, and V. Bernardo. Optimal location of fast charging stations in barcelona: A flow-capturing approach. In *10th International Conference on the European Energy Market (EEM)*, pages 1–6, 2013.

[16] S. Dempe. *Foundations of Bilevel Programming.* Springer, Berlin, 2002.

[17] S. Dempe. *Bilevel Optimization: Theory, Algorithms and Applications.* Springer, Cham, 2018.

[18] Eurostat. Geostat population grid dataset. `https://ec.europa.eu/eurostat/web/gisco/geodata/population-distribution/population-grids`, 2023. Accessed: 2025-02-10.

[19] J. M. Fortuny-Amat and B. J. McCarl. A representation and economic interpretation of a system of complementary inequalities. *Operations Research*, 29(4):811–816, 1981.

[20] L. Franceschi, P. Frasconi, S. Salzo, R. Grazzi, and M. Pontil. Bilevel programming for hyperparameter optimization and meta-learning. *Proceedings of the 35th International Conference on Machine Learning (ICML)*, pages 1225–1234, 2018.

[21] I. P. Gasbarro, A. Pizzuti, and O. E. Karaşan. Bilevel design and pricing of ev charging stations with deviation-flow. In *Proc. Int. Conf. on Optimization and Decision Science (ODS 2025)*, 2025.

[22] F. Guo, J. Yang, and J. Lu. The battery charging station location problem: Impact of users' range anxiety and distance convenience. *Transportation Research Part E: Logistics and Transportation Review*, 114:1–18, 2018.

[23] Y. Honma and M. Kuby. Node-based vs. path-based location models for urban hydrogen refueling stations: Comparing convenience and coverage abilities. *International Journal of Hydrogen Energy*, International Journal of Hydrogen Energy, 2017.

[24] Y. Jiang, Y. Zhang, C. Zhang, and J. Fan. The capacitated deviation-flow fueling location model for siting battery-charging stations. In *Proceedings of the 12th COTA International Conference of Transportation Professionals (CICTP 2012)*, pages 2771–2778, 2012.

[25] M. Kendall. A new measure of rank correlation. *Biometrika*, 30(1-2):81–93, 1938.

[26] J. Kim and M. Kuby. The deviation-flow refueling location model for optimizing a network of refueling stations. *International Journal of Hydrogen Energy*, 37(6):5406–5420, 2012.

[27] J.-G. Kim and M. Kuby. A network transformation heuristic approach for the deviation flow refueling location model. *Computers & Operations Research*, 40(5):1122–1131, 2013.

[28] T. Kleinert, M. Köppe, M. Lübbecke, and I. Nowak. A survey on mixed-integer bilevel optimization. *Mathematical Programming*, 188(1):251–301, 2021.

[29] C. Kong, R. Jovanovic, I. Bayram, and M. Devetsikiotis. A hierarchical optimization model for a network of electric vehicle charging stations. *Energies*, 2016.

[30] M. Kuby and S. Lim. The flow-refueling location problem for alternative-fuel vehicles. *Socio-Economic Planning Sciences*, 39(2):125–145, 2005.

[31] Kınay, F. Gzara, and S. Alumur. Full cover charging station location problem with routing. *Transportation Research Part B: Methodological*, 144:1–22, 2021.

[32] Kınay, F. Gzara, and S. Alumur. Charging station location and sizing for electric vehicles under congestion. *Transportation Science*, 57:343–354, 2023.

[33] M. Labbé, P. Marcotte, and G. Savard. A bilevel model of taxation and its application to optimal highway pricing. *Management Science*, 44:1608–1622, 1998.

[34] C. Lee and J. Han. Benders-and-price approach for electric vehicle charging station location problem under probabilistic travel range. *Transportation Research Part B: Methodological*, 106:130–152, 2017.

[35] S. Lim and M. Kuby. Heuristic algorithms for siting alternative-fuel stations using the flow-refueling location model. *European Journal of Operational Research*, 204(1):51–61, 2010.

[36] Z.-Q. Luo, J.-S. Pang, and D. Ralph. *Mathematical Programs with Equilibrium Constraints.* Cambridge University Press, Cambridge, 1996.

[37] M. Miralinaghi, B. Keskin, Y. Lou, and A. Roshandeh. Capacitated refueling station location problem with traffic deviations over multiple time periods. *Networks and Spatial Economics*, 17, 2017.

[38] S. A. MirHassani and R. Ebrazi. A flexible reformulation of the refueling-station location problem. *Transportation Science*, 47(4):617–628, 2013.

[39] OpenStreetMap, 2025. Accessed: 2025-05-27.

[40] J. V. Outrata, M. Kočvara, and J. Zowe. *Nonsmooth Approach to Optimization Problems with Equilibrium Constraints*. Kluwer Academic Publishers, Dordrecht, 1998.

[41] K. Pearson. Notes on regression and inheritance in the case of two parents. *Proceedings of the Royal Society of London*, 58:240–242, 1895.

[42] F. Pedregosa. Hyperparameter optimization with approximate gradient. *Proceedings of the 33rd International Conference on Machine Learning (ICML)*, pages 737–746, 2016.

[43] Regione Lombardia. Matrice od2016 – passeggeri – scenario 2020. `https://www.dati.lombardia.it/Mobilit-e-trasporti/Matrice-OD2016-Passeggeri-Scenario-2020/hyqr-mpe2`, 2020. Open data, accessed 28 November 2025.

[44] R. Riemann, D. Wang, and F. Busch. Optimal location of wireless charging facilities for electric vehicles: Flow capturing location model with stochastic user equilibrium. *Transportation Research Part C: Emerging Technologies*, 2015.

[45] R. Sa'adati, M. Jafari-Nokandi, and J. Saebi. Allocation of ress and pev fast-charging stations on coupled transportation and distribution networks. *Energy Reports*, 7:5962–5977, 2021.

[46] S. Shapiro and M. Wilk. An analysis of variance test for normality (complete samples). *Biometrika*, 52(3-4):591–611, 1965.

[47] D. Simchi-Levi and O. Berman. A heuristic algorithm for the traveling salesman location problem on networks. *Operations Research*, 36(3):478–484, 1988.

[48] F. Simini, M. C. González, A. Maritan, and A.-L. Barabási. A universal model for mobility and migration patterns. *Nature*, 484(7392):96–100, 2012.

[49] A. Sinha, M. Malo, and K. Deb. Review of bilevel optimization and applications in machine learning. *IEEE Transactions on Evolutionary Computation*, 22(2):256–270, 2017.

[50] C. Spearman. The proof and measurement of association between two things. *American Journal of Psychology*, 15:72–101, 1904.

[51] J. Tan and W.-H. Lin. A stochastic flow capturing location and allocation model for siting electric vehicle charging stations. *2014 17th IEEE International Conference on Intelligent Transportation Systems, ITSC 2014*, 2014.

[52] TomTom International BV. Tomtom traffic data – real-time and historical traffic information. `https://www.tomtom.com/traffic-data/`, 2025. Accessed June 5, 2025.

[53] C. Upchurch, M. Kuby, and S. Lim. A model for location of capacitated alternative-fuel stations. *Geographical Analysis*, 41(1):85–106, 2009.

[54] H. von Stackelberg. *Marktform und Gleichgewicht.* Springer, Vienna, 1934.

[55] F. Wu and R. Sioshansi. A stochastic flow-capturing model to optimize the location of fast-charging stations with uncertain electric vehicle flows. *Transportation Research Part D: Transport and Environment*, 53:354–376, 2017.

[56] M. Xu and Q. Meng. Optimal deployment of charging stations considering path deviation and nonlinear elastic demand. *Transportation Research Part B: Methodological*, 135:120–142, 2020.

[57] B. Yıldız, O. Arslan, and O. Karaşan. A branch and price approach for routing and refueling station location model. *European Journal of Operational Research*, 248:815–826, 2016.

[58] B. Yıldız, E. Olcaytu, and A. Şen. The urban recharging infrastructure design problem with stochastic demands and capacitated charging stations. *Transportation Research Part B: Methodological*, 119:22–44, 2019.

[59] E. Zavvos, E. Gerding, and M. Brede. A comprehensive game-theoretic model for electric vehicle charging station competition. *IEEE Transactions on Intelligent Transportation Systems*, 23(8):12239–12250, 2022.

[60] B. Zeng and Y. An. Solving mixed-integer bilevel programs by a hybrid reformulation–decomposition method. *Operations Research*, 62(3):665–678, 2014.

[61] H. Zhang, S. J. Moura, Z. Hu, and Y. Song. Pev fast-charging station siting and sizing on coupled transportation and power networks. *IEEE Transactions on Smart Grid*, PP:1–1, 2017.

Appendix A

Proofs

A.1 Variance Reduction by Averaging and the Role of Decorrelation

This section provides formal results and proofs showing how averaging multiple predictors reduces variance, why decorrelation between base learners is essential (as in Random Forests), and how weighting can further minimize ensemble variance. Throughout, we consider a regression setting with a (fixed) input $\mathbf{x}$ and random predictors $h_b(\mathbf{x})$, $b = 1, \ldots, B$, built from the same training distribution but perturbed via resampling, random feature selection, or stochastic optimization. To simplify notation, we suppress the dependence on $\mathbf{x}$ and write h_b.

A.1.1 Setup and Notation

Let $\mathbf{h} = (h_1, \ldots, h_B)^\top$ be the vector of base learners. Denote their mean $\mu_b = \mathbb{E}[h_b]$ and variance $\sigma_b^2 = \mathrm{Var}(h_b)$, and the covariance matrix $\Sigma \in \mathbb{R}^{B\times B}$ with entries $\Sigma_{ij} = \mathrm{Cov}(h_i, h_j)$. Consider the (possibly weighted) ensemble

$$\hat{f}_w \;=\; \sum_{b=1}^{B} w_b\, h_b, \qquad \text{with} \quad \sum_{b=1}^{B} w_b = 1, \tag{A.1}$$

and the special case of uniform averaging (bagging),

$$\hat{f}_{\mathrm{avg}} \;=\; \frac{1}{B}\sum_{b=1}^{B} h_b. \tag{A.2}$$

A.1.2 Variance of a General Weighted Ensemble

[Ensemble Variance] For weights w summing to 1, the ensemble variance is

$$\mathrm{Var}(\hat{f}_w) \;=\; w^\top \Sigma\, w. \tag{A.3}$$

By linearity and bilinearity of covariance,

$$\mathrm{Var}\left(\sum_b w_b h_b\right) = \sum_{i,j} w_i w_j \,\mathrm{Cov}(h_i, h_j) = w^\top \Sigma w.$$

A.1.3 Uniform Averaging: General and Symmetric Cases

[Variance under Uniform Averaging] For $\hat{f}_{\text{avg}}$ as in (A.2),

$$\mathrm{Var}(\hat{f}_{\text{avg}}) \;=\; \frac{1}{B^2}\left(\sum_{b=1}^{B} \sigma_b^2 \;+\; 2 \sum_{1\le i<j\le B} \mathrm{Cov}(h_i, h_j)\right). \tag{A.4}$$

Apply Proposition A.1.2 with $w_b = 1/B$ and expand.

[Equal Variance and Equal Pairwise Correlation] Suppose $\mathrm{Var}(h_b) = \sigma^2$ for all b and $\mathrm{Corr}(h_i, h_j) = \rho$ for all $i \neq j$ (so $\mathrm{Cov}(h_i, h_j) = \rho\sigma^2$). Then

$$\mathrm{Var}(\hat{f}_{\text{avg}}) \;=\; \frac{\sigma^2}{B}\Big(1 + (B-1)\rho\Big). \tag{A.5}$$

Insert $\sigma_b^2 = \sigma^2$ and $\mathrm{Cov}(h_i, h_j) = \rho\sigma^2$ into (A.4):

$$\mathrm{Var}(\hat{f}_{\text{avg}}) = \frac{1}{B^2}\left(B\sigma^2 + 2\binom{B}{2}\rho\sigma^2\right) = \frac{1}{B^2}\left(B\sigma^2 + B(B-1)\rho\sigma^2\right) = \frac{\sigma^2}{B}\big(1 + (B-1)\rho\big).$$

Consequences.

- If $\rho = 0$ (uncorrelated base learners), then $\mathrm{Var}(\hat{f}_{\text{avg}}) = \sigma^2/B$, i.e., variance decays at the canonical rate $1/B$.
- If $0 < \rho < 1$, then as $B \to \infty$,

  $$\mathrm{Var}(\hat{f}_{\text{avg}}) \;\to\; \rho\,\sigma^2, \tag{A.6}$$

  so variance cannot be reduced below $\rho\sigma^2$ by mere averaging; decorrelation is essential.
- If $\rho = 1$ (perfectly correlated), then $\mathrm{Var}(\hat{f}_{\text{avg}}) = \sigma^2$: averaging yields *no* variance reduction.

These points explain the Random Forest design: bootstrap sampling and random feature subspaces aim to *reduce* ρ, thereby lowering the variance floor.

A.1.4 Ensemble Improves upon a Single Base Learner

[Variance Reduction Criterion] Under the symmetric assumptions of Proposition A.1.3, for any $B \ge 2$ and any $\rho < 1$,

$$\mathrm{Var}(\hat{f}_{\text{avg}}) \;<\; \sigma^2 \quad\Longleftrightarrow\quad \frac{1}{B}\big(1 + (B-1)\rho\big) \;<\; 1 \quad\Longleftrightarrow\quad \rho \;<\; 1. \tag{A.7}$$

Immediate from (A.5).

A.1.5 Bias–Variance Decomposition for Averaging

Let $f^\star$ denote the (fixed) target function value at $\mathbf{x}$. The squared error decomposes as

$$\mathbb{E}\Big[(\hat{f}_{\text{avg}} - f^\star)^2\Big] = \underbrace{\big(\mathbb{E}[\hat{f}_{\text{avg}}] - f^\star\big)^2}_{\text{Bias}^2} + \underbrace{\text{Var}(\hat{f}_{\text{avg}})}_{\text{Variance}} + \underbrace{\sigma_\epsilon^2}_{\text{Irreducible noise}}, \tag{A.8}$$

where σ_ϵ^2 is the observation noise variance (if present). If each h_b is (approximately) unbiased for $f^\star$, then averaging preserves low bias while reducing variance according to (A.4)–(A.5). This is the canonical explanation for Random Forests: variance drops substantially while bias stays roughly unchanged relative to a deep tree.

A.1.6 Unequal Variances and Correlations: A Useful Bound

Define the average variance $\bar{\sigma}^2 = \frac{1}{B}\sum_{b=1}^{B}\sigma_b^2$ and the average pairwise correlation

$$\bar{\rho} \;=\; \frac{2}{B(B-1)} \sum_{1\le i<j\le B} \frac{\text{Cov}(h_i, h_j)}{\sigma_i\sigma_j}. \tag{A.9}$$

Then, by the AM–GM inequality and Cauchy–Schwarz,

$$\text{Var}(\hat{f}_{\text{avg}}) = \frac{1}{B^2}\left(\sum_{b=1}^{B}\sigma_b^2 + 2\sum_{i<j}\text{Cov}(h_i, h_j)\right) \;\le\; \frac{\bar{\sigma}^2}{B} \;+\; \frac{B-1}{B}\,\bar{\rho}\,\bar{\sigma}^2. \tag{A.10}$$

Thus the same qualitative message holds: variance contracts like $1/B$ when $\bar{\rho} \approx 0$, but is lower-bounded by $\bar{\rho}\,\bar{\sigma}^2$ in the large-B limit.

A.1.7 Optimal Weights under Unbiasedness

Suppose $\mathbb{E}[h_b] = f^\star$ for each b (unbiasedness) and we seek weights w summing to 1 that minimize ensemble variance.

[Minimum-Variance Unbiased Ensemble] Among all w with $\mathbf{1}^\top w = 1$, the variance $w^\top \Sigma w$ is minimized by

$$w^\star \;=\; \frac{\Sigma^{-1}\mathbf{1}}{\mathbf{1}^\top\Sigma^{-1}\mathbf{1}}, \qquad \text{Var}(\hat{f}_{w^\star}) = \frac{1}{\mathbf{1}^\top\Sigma^{-1}\mathbf{1}}. \tag{A.11}$$

Use Lagrange multipliers for $\min_w w^\top\Sigma w$ subject to $\mathbf{1}^\top w = 1$:

$$\mathcal{L}(w, \lambda) = w^\top\Sigma w + \lambda(\mathbf{1}^\top w - 1).$$

Setting $\nabla_w \mathcal{L} = 0$ yields $2\Sigma w + \lambda\mathbf{1} = 0$, hence $w = -\frac{\lambda}{2}\Sigma^{-1}\mathbf{1}$. Enforcing $\mathbf{1}^\top w = 1$ gives the stated $w^\star$ and value.

Special Cases.

- If Σ is diagonal (uncorrelated base learners), then $w_b^\star \propto 1/\sigma_b^2$ (inverse-variance weighting).
- If, in addition, $\sigma_b^2 \equiv \sigma^2$, then $w^\star$ is uniform: $w_b^\star = 1/B$.

A.1.8 From Bagging to Random Forests: Why Random Subspaces Help

Proposition A.1.3 shows that the limiting variance floor is $\rho\sigma^2$. Random Forests reduce ρ by (i) training each tree on a bootstrap sample (data perturbation) and (ii) restricting each split to a random subset of features (model perturbation). Both mechanisms decrease cross-tree covariance by forcing different trees to rely on distinct samples and distinct predictive cues, thereby tightening (A.10) and lowering the variance floor.

A.1.9 A Simple Bias–Variance Comparison: Single Tree vs. Random Forest

Let a single deep tree have bias B and variance σ^2. For a Random Forest of B such trees with pairwise correlation ρ, assuming similar bias across trees, the ensemble bias remains approximately B while the variance becomes $\frac{\sigma^2}{B}\big(1 + (B-1)\rho\big)$. Hence the Random Forest improves mean squared error (MSE) over a single tree whenever

$$\frac{\sigma^2}{B}\big(1 + (B-1)\rho\big) \;+\; \mathsf{B}^2 \;<\; \sigma^2 \;+\; \mathsf{B}^2 \quad\Longleftrightarrow\quad \rho < 1, \tag{A.12}$$

i.e., as long as trees are not perfectly correlated.

A.1.10 Takeaways

- Averaging reduces variance in proportion to the number of base learners when they are weakly correlated.
- The large-B variance limit equals the average correlation times the average variance; decorrelation is therefore paramount.
- Optimal (minimum-variance) weights depend on Σ^{-1}; uniform weights are optimal under homoscedastic, uncorrelated base learners.

Tesi redatta con il cofinanziamento dell'Unione europea-*NextGeneration EU*
Piano Nazionale di Ripresa e Resilienza Missione 4, Componente 2 –
Investimento 3.3. del PNRR – "Dottorati innovativi che rispondono ai fabbisogni di innovazione delle imprese" – CUP E11I22000200001

Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA
UNIVERSITÀ DEGLI STUDI DELL'AQUILA