

Determining the Location of Charging Station for a One-Way Electric Car-Sharing System Under Demand Uncertainty

Zhaoyu Yan

2024



Determining the Location of Charging Station for a One-Way Electric Car-Sharing System Under Demand Uncertainty

Zhaoyu Yan

October 2024

Master thesis submitted to Delft University of Technology in partial fulfilment of the requirements for the degree

Master of Science in Civil Engineering – Traffic and Transport Engineering

To be defended in public on October 31st, 2024

Graduation Committee

Chairperson, TU Delft

First Supervisor, TU Delft

Second Supervisor, TU Delft

Prof. Dr. Oded Cats

Dr. Jie Gao

Dr. Yousef Maknoon

Acknowledgments

Time flies, and in the blink of an eye, my two years of studying abroad are coming to an end with the completion of this thesis. Looking back over these two years, countless memories flood my mind: the initial anxiety of arriving in a foreign country, the gradual process of figuring out how to adapt to the academic demands at TU Delft, the encouragement found in studying for exams with my classmates, and the anxiety of searching for the career path that suits me best. These fragments of memories come together to form a rich and unforgettable chapter in my life.

I am deeply grateful to TU Delft for providing such a valuable platform. Here, I was fortunate to learn a wide range of technical knowledge, hone my problem-solving thinking ability, and discover new ways of exploring this world.

I would like to express my heartfelt thanks to my committee for their invaluable guidance throughout this journey. Your insights and suggestions have been of immense benefit to me. First, I want to thank the chair of my committee, Prof. Oded, for your significant contribution to shaping the direction of my thesis. I vividly remember how, at the start, you took the time to carefully revise my early-stage conference document, offering detailed feedback despite your busy schedule. You not only corrected technical aspects but also pointed out grammatical errors, which deeply moved me. Your thoughtful guidance in our meetings has been invaluable, and I am sincerely thankful for your support. Next, I would like to thank my first supervisor, Dr. Jie, who is also a good friend. You have provided immense guidance throughout my entire thesis process. You were the one who first introduced me to the world of operations research, guiding me from foundational concepts to model building, and then step by step into deeper research. I must say, the combination of operations research and deep learning is truly fascinating, and I have gained a great deal of technical knowledge from you. Beyond academic guidance, you have also frequently offered me advice and encouragement about life. You often shared your personal growth experiences and offered genuine suggestions from your own journey, and I have learned so much from you. For all of this, I am deeply grateful. Lastly, I want to thank my second supervisor, Dr. Yousef. Your approach to optimization has been incredibly enlightening for me. Every meeting with you has been a valuable learning experience, and I have constantly admired your way of thinking. Your guidance has been like a ladder, helping me to continuously improve, and for that, I am truly thankful.

I would also like to thank my girlfriend and my close friends who have been by my side throughout this journey. Your encouragement and companionship have made me feel less alone in a foreign land and have given me the strength to keep moving forward.

Finally, I want to give my deepest thanks to my parents. It is your unwavering support that has made it possible for me to come to the Netherlands and fulfill my dream of studying abroad. Despite the thousands of kilometers between us, every phone call made me feel the warmth of home. Thank you for everything.

Abstract

This thesis explores the optimization of charging station location within a one-way electric car-sharing system, addressing the inherent challenges of demand uncertainty. We introduce a novel deep learning-based stochastic programming framework (LMSP Framework) to tackle this issue. This framework integrates two key components:

1. A deep learning model (LSTM-MLP-MDN), composed of Long Short-Term Memory (LSTM), Multilayer Perceptron (MLP), and Mixture Density Network (MDN) architectures, which predicts the probability distribution of traffic demand using historical data.
2. A two-stage stochastic programming model, designed to strategically optimize the locations of charging stations and the initial number of vehicles at each station under demand uncertainty.

The primary goal of this framework is to improve the profitability and operational efficiency of the car-sharing system by optimizing both the location and capacity of charging stations, effectively solving the Charging Station Location Problem (CSLP).

We validate the effectiveness, adaptability, and feasibility of our framework through a comprehensive case study in Manhattan, utilizing historical traffic data to ensure the reliability of the deployment plan. Our results demonstrate that integrating deep learning techniques with stochastic programming significantly enhances both the accuracy of demand forecasting and the consistency of resource allocation in the optimization process for charging station locations. Specifically, the LMSP Framework achieves higher operational efficiency in the short term, as evidenced by superior metrics like Demand Satisfaction Ratio (DSR) and Charging Station Utilization Rate (CSU) compared to traditional methods. This ensures more balanced and efficient resource allocation across different demand scenarios. However, traditional approaches tend to perform better in short-term financial indicators, such as profit and Return on Investment (ROI), as they employ more aggressive resource allocation strategies based on higher demand forecasts. Despite this, the LMSP Framework's focus on operational efficiency positions it as a more viable option for long-term profitability and user satisfaction, offering a sustainable solution for urban mobility.

Furthermore, our analysis provides valuable recommendations for future charging station deployments. These findings have important implications for the planning and operation of electric car-sharing systems, potentially contributing to more sustainable and efficient urban mobility solutions for all stakeholders involved.

Keywords: Demand Uncertainty; Charging Station Location Problem; One-way Car-sharing; Deep learning; Stochastic Programming

Contents

1. Introduction	2
1.1. Background Description	2
1.2. Objectives and Research Scope	4
1.2.1. Research Objectives	4
1.2.2. Research Scope	5
1.3. Research Questions	5
1.4. Thesis Outline	6
2. Literature review	8
2.1. Charging Station Location Problem in One-Way Car-Sharing Systems	8
2.1.1. Traditional Methods	8
2.1.2. Robust Optimization (RO)	9
2.1.3. Stochastic Programming	10
2.1.4. Key Findings and Implications	12
2.2. Methods for Traffic Demand Prediction	13
2.2.1. Classical Methods	13
2.2.2. Deep learning Methods	14
2.3. Research Gap and Main Contributions	20
3. Problem Description and Formulation	24
3.1. Problem Description	24
3.2. Two-Stage Stochastic Programming Model	25
3.2.1. Assumptions	25
3.2.2. Notations and Definitions	25
3.2.3. Time-Dependent Integer Linear Program	29
4. Methodology	33
4.1. Overview of the LMSP Framework	33
4.2. Deep Learning Model for Traffic Demand Forecasting	34
4.2.1. LSTM-MLP-MDN Architecture	34
4.2.2. Multilayer Perceptron (MLP)	34
4.2.3. Long Short-Term Memory (LSTM)	35
4.2.4. Mixture Density Network (MDN)	36
5. Case Study	38
5.1. Dataset and Preprocessing	38
5.1.1. New York City (NYC) taxi dataset	38
5.1.2. Contextual Data of Manhattan	42
5.2. Traffic Demand Prediction	43
5.2.1. Experiment Setup	43
5.2.2. Evaluation Metric	45
5.2.3. Comparison With State-of-the-Art	47

Contents

5.2.4. Results Analysis	48
5.3. Charging Station Deployment Plan	50
5.3.1. Determining Critical Time Periods	50
5.3.2. Assigning Predicted Traffic Demand to Specific Routes	51
5.3.3. Traffic Demand Level Scenario Sampling	51
5.3.4. Optimal Charging Station Deployment Results	54
5.3.5. The Added Value of Deep Learning in Stochastic Programming for Demand Prediction	54
5.4. Comprehensive Analysis for Charging Station Location Optimization	57
5.4.1. Overall Demand Patterns	57
5.4.2. Spatial Demand Patterns	58
5.4.3. Evaluation of Deployment Plan	60
5.4.4. Deployment Plan Recommendations	71
6. Discussion	76
6.1. Discussion of Results	76
6.1.1. Effectiveness of the LMSP Framework	76
6.1.2. Deployment Strategy Analysis	77
6.2. Discussion of Assumptions and Limitations	77
6.2.1. Assumptions	77
6.2.2. Limitations	79
6.3. Generalisability of the model and results	80
6.3.1. Applicability to Urban Areas	80
6.3.2. Scalability to Other Cities	80
6.3.3. Data Availability and Model Training	81
6.3.4. Limitations of Generalisation to Short-term or Real-time Adjustments	81
7. Conclusion	83
7.1. Answers to the Research Questions	83
7.1.1. Main Research Question	83
7.1.2. Sub Research Questions	84
7.2. Future Research Outlook	87
A. Deployment Plan over Five Time Periods	90
B. Trip Data at Different Locations over five specific time periods	96
C. Fundamentals of Deep Learning Model	98
C.1. LSTM unit	98
C.2. MDN	99
C.3. Activation Function	100
D. Results of Comparison between LMSP Framework and Traditional-based SP	101

Acronyms

EV	Electric Vehicle
CSLP	Charging Station Location Problem
RO	Robust Optimization
SP	Stochastic Programming
MLP	Multilayer Perceptron
LSTM	Long Short-Term Memory
MDN	Mixture Density Network
NYC	New York City
LMSP	LSTM-MLP-MDN Stochastic Programming
KPI	Key Performance Indicator
HFA	Historical Frequency-based Approach
GDA	Gaussian Distribution Approach

List of Figures

1.1. One-way Car-sharing System	2
1.2. Electric Car-Sharing Service in Los Angeles	3
2.1. Multilayer Perceptron (MLP)	15
2.2. Research gap	21
3.1. Time-expanded Location Graph	29
4.1. Methodological Framework	33
4.2. LSTM-MLP-MDN Architecture	35
4.3. The mechanism of LSTM	36
5.1. Taxis of New York City	38
5.2. Heatmap of Manhattan Traffic Demand from the Perspectives of Departure and Arrival	39
5.3. 25 Potential Locations of Charging Stations	41
5.4. Sliding window cross-validation	44
5.5. Training and Validation Loss over Five Windows	44
5.6. Total Traffic Demand on January 1st by Year (2015-2020), Predicted for January 1st, 2020	49
5.7. Predicted Total traffic demand in Manhattan on January 1, 2020	50
5.8. Performance Comparison of Different Methods Over Time Using KPIs	56
5.9. Trips by starting region and hour	58
5.10. Spatial Distribution at 4h	62
5.11. Spatial Distribution at 9h	64
5.12. Spatial Distribution at 15h	66
5.13. Spatial Distribution at 18h	68
5.14. Spatial Distribution at 21h	70
C.1. Mixture Density Network structure [Petrov and Repin, 2020]	99
C.2. ReLU Activation Function	100
C.3. Softplus Activation Function	100

List of Tables

2.1. Comparison of three traditional machine learning methods	14
2.2. A Review of Deep Learning Methods in Traffic Demand Prediction Research	19
3.1. Description of different types of arcs	29
5.1. 25 Potential Locations of Charging Stations	42
5.2. Hourly Traffic Demand Hyperparameters	43
5.3. Comparison of the model results for hourly prediction	48
5.4. Probabilistic Distribution of Predicted Traffic Demand Over Five Time Periods	52
5.5. Model Parameters	53
5.6. Gurobi Setup	53
5.7. Summary of Station Metrics at Different Times	71
5.8. Deployment Plan Recommendations Table	74
A.1. Deployment Plan at 4h	91
A.2. Deployment Plan at 9h	92
A.3. Deployment Plan at 15h	93
A.4. Deployment Plan at 18h	94
A.5. Deployment Plan at 21h	95
B.1. Trip Data at Different Locations	97
D.1. Performance at 4h	101
D.2. Performance at 9h	101
D.3. Performance at 15h	101
D.4. Performance at 18h	101
D.5. Performance at 21h	101



1

Introduction

1. Introduction

1.1. Background Description

The widespread ownership of private vehicles presents a complex challenge, creating multiple interconnected issues. For instance, in the European Union (EU), road transport accounts for about 20% of the EU's total greenhouse gas emissions. Passenger cars are a major polluter, responsible for 61% of total CO₂ emissions from EU road transport. Most of these vehicles run on conventional fuels, significantly contributing to environmental pollution and global warming, with urban areas experiencing severe air quality degradation.

The transition to electric vehicles (EVs) serves as a practical approach to reducing these environmental concerns. Unlike conventional vehicles that run on fossil fuels, EVs produce zero tailpipe emissions, significantly reducing air pollution and greenhouse gas emissions. According to the International Energy Agency, the CO₂ emissions from EVs can be over 50% lower than those from internal combustion engine vehicles, even when accounting for the emissions from electricity generation [International Energy Agency, 2022]. Furthermore, the development of smart grid technologies and energy storage solutions enhances the sustainability of electric vehicles. The United Nations recognizes the potential of electric mobility in achieving global environmental goals. Through the Global Electric Mobility Programme, the UN supports over 50 low- and middle-income countries in transitioning from fossil fuels to electric vehicles [United Nations Environment Programme, 2024]. The shift to electric vehicles is critical in addressing these environmental concerns, reducing our reliance on fossil fuels, and mitigating climate change impacts.




Figure 1.1.: One-way Car-sharing System

Car-sharing offers a sustainable alternative to private car ownership by providing flexible, short-term vehicle access, which can reduce urban congestion, parking demand, and environmental impacts (see [Ke et al., 2019; Baptista et al., 2014; Glotz-Richter, 2016]). Users

1. Introduction



Figure 1.2.: Electric Car-Sharing Service in Los Angeles

can locate and reserve vehicles through mobile apps, paying only for the duration of use. Traditionally, car-sharing models can be divided into three types: two-way station-based model (see [Strobel and Pruckner, 2022]), one-way station-based model (see [Shaheen et al., 2015; Febbraro et al., 2019; Nourinejad and Roorda, 2015]), and free-floating model. In one-way systems, users can pick up a car from one location and drop it off at another. In two-way systems, users must return the car to the same location from which they picked it up [Nourinejad and Roorda, 2015]. The free-floating model differs significantly from the station-based models by allowing users to pick up and drop off vehicles at any legal parking spot within a predefined operational area (see [Firnkorn and Müller, 2011; Formentin et al., 2015; Cocca et al., 2019]). This research focuses on a one-way system due to its flexibility and convenience for users. For example, consider a commuter who uses a one-way car-sharing service to drive from a suburban home to a downtown office in the morning, then parks the vehicle at a charging station for other users, and later takes another car to return home in the evening (see Figure 1.1). This flexibility not only saves time but also reduces the need for parking spaces at the user’s destination. However, one-way systems require more vehicles to be available at different locations, leading to operational challenges in balancing the distribution of cars across the service area.

Combining electric vehicles with car-sharing services can amplify the benefits of both technologies, extending beyond environmental impacts. It provides users with the flexibility and convenience of traditional car-sharing while also reducing emissions and dependency on fossil fuels [Clewlow, 2016]. Major cities worldwide have adopted electric car-sharing programs to promote sustainable transportation. For example, in the United States, cities such as Los Angeles and Denver have embraced electric car-sharing with a focus on equity. In China, Didi Chuxing, the leading ride-hailing company, has launched an electric car-sharing service in partnership with major automakers like Renault-Nissan-Mitsubishi. This program includes a large fleet of electric vehicles and aims to significantly expand the availability of electric vehicles across Chinese cities. The electric car-sharing market is expected to thrive, driven by a combination of stringent government regulations, increasing environmental awareness, and technological advancements.

However, one of the significant challenges in the operation of one-way electric car-sharing systems is demand uncertainty. Predicting when and where users will need cars and charg-

ing stations can be complex due to varying factors such as time of day, season, weather conditions, traffic conditions, and local events. This unpredictability makes it difficult to determine optimal locations for charging stations [Giménez-Gaydou et al., 2016]. The operational consequences of this uncertainty are significant: overestimating demand can lead to resource oversupply and inflated costs, while underestimating demand may cause service disruptions and lead to dissatisfied customers. Although traditional approaches have attempted to address this issue, they often fail to capture the complexity and dynamic nature of traffic demand patterns in urban environments. This highlights the pressing need for more advanced, data-driven approaches that can more effectively handle demand fluctuations and optimize the location of charging stations in electric car-sharing systems.

1.2. Objectives and Research Scope

1.2.1. Research Objectives

This research introduces a novel framework, termed the LMSP Framework, which integrates advanced deep learning techniques—Long Short-Term Memory (LSTM), Multilayer Perceptron (MLP), and Mixture Density Networks (MDN)—with stochastic programming. This innovative approach tackles the critical challenge of optimizing charging station locations and the allocation of initial vehicle resources under demand uncertainty. By applying the predictive power of machine learning and the robustness of stochastic optimization, we aim to create a more adaptive and effective solution to the Charging Station Location Problem (CSLP). The primary objectives are outlined as follows:

- **Develop and optimize the LMSP Framework:** Design and optimize the LMSP framework to address the Charging Station Location Problem (CSLP). This framework will integrate an LSTM-MLP-MDN deep learning model to predict traffic demand probability distributions and incorporate these predictions into a two-stage stochastic programming model to optimize the locations of charging stations and initial vehicle resources.
- **Predict a more accurate traffic demand probability distribution:** Utilize deep learning techniques to predict a more precise traffic demand probability distribution based on historical traffic data. This will generate a realistic set of traffic demand scenarios along with their corresponding probabilities, which will be incorporated into a stochastic programming model to more effectively manage demand uncertainty and provide a solid foundation for charging infrastructure deployment.
- **Validate the effectiveness, adaptability, and feasibility of the LMSP Framework in similar urban settings:** Through a comprehensive case study in a real urban environment, validate the proposed LMSP Framework’s effectiveness, adaptability in dealing with demand uncertainty, and feasibility in similar high-density urban settings.
- **Support sustainable urban transportation through infrastructure planning:** This research is committed to providing a method that enhances the planning of electric car-sharing infrastructure to support sustainable urban transportation goals. By optimizing resource allocation and infrastructure utilization, we aim to alleviate urban congestion, reduce environmental pollution, and promote more sustainable urban development.

1.2.2. Research Scope

The scope of the research primarily includes the following aspects:

- **One-way Car-sharing System:** This research focuses on addressing the CSLP in a one-way car-sharing system, and may not be directly applicable to two-way or free-floating car-sharing models.
- **Key Decision:** This research focuses on charging station location and the allocation of vehicle resources. It does not include operational decisions such as rebalancing [Guo and Kang, 2022], charging scheduling [Xie et al., 2020], or price strategies [Wang et al., 2021].
- **Demand Uncertainty:** Demand uncertainty is the only uncertainty parameter considered in this research, while other potential uncertainties, such as driving range uncertainty and vehicle arrival time uncertainty (see 2.1.3), are not taken into account. Furthermore, dynamic factors such as economic conditions and policy changes [Zhang et al., 2018][Guo et al., 2021][Wang et al., 2019b], which may affect the model's real-world performance, are also excluded.
- **Demand Pattern Temporal Focus:** This research focuses on short-term demand patterns. Specifically, we analyze hourly fluctuations in traffic demand.
- **Applicability to urban environments:** This research selects Manhattan as the case city, as it represents a typical high-density urban area characterized by highly representative traffic patterns and demand features. Therefore, the research is mainly applicable to similar high-density, high-traffic demand urban areas, and its relevance to low-density cities or rural regions may be limited.

By establishing these clear research scopes, we aim to provide a focused and in-depth analysis of the CSLP within the context of one-way car-sharing systems in high-density urban environments, while acknowledging the potential for future research to address more broadly issues related to EV charging station deployment.

1.3. Research Questions

Main Research Question

How can a novel deep learning-based stochastic programming framework (LMSP Framework) be developed and applied to optimize the Charging Station Location Problem (CSLP) in a one-way electric car-sharing system under conditions of traffic demand uncertainty?

Sub Research Questions

1. How does the integration of stochastic programming help in accommodating demand uncertainty when optimizing the location and initial number of cars at charging stations?
2. How can advanced deep learning techniques, particularly the LSTM-MLP-MDN model, be employed to accurately forecast the probability distribution of traffic demand in a one-way electric car-sharing system?

3. How does the integration of advanced demand prediction techniques, such as LSTM-MLP-MDN, enhance the performance of stochastic programming models addressing the CSLP compared to traditional forecasting methods, and what added value does it bring?
4. What insights can be drawn from the application of the proposed LMSP Framework in a real-world case study (Manhattan)?
5. How do the results of the case study validate the effectiveness, adaptability, and feasibility of the proposed LMSP Framework?
6. How can the findings from this research enhance the planning and operation of electric car-sharing systems, and what are the potential implications for sustainable urban mobility?

1.4. Thesis Outline

The remainder of this paper is organized as follows:

Chapter 2 presents a review of related works in the literature. This review establishes the theoretical foundation for our research and identifies the gaps in current knowledge that our research aims to address.

Chapter 3 provides a detailed description of the research problem and formulates the two-stage stochastic programming mathematical model related to the Charging Station Location Problem (CSLP).

Chapter 4 introduces the LMSP Framework and provides a detailed explanation of the deep learning model used.

Chapter 5 presents a case study in Manhattan to validate the LMSP Framework.

Chapter 6 provides a discussion on the results, including the effectiveness of the LMSP Framework and the analysis of deployment strategy. It then addresses the assumptions and limitations of this research.

Chapter 7 concludes by summarizing the key findings, addressing all the research questions, and offering insights into potential directions for future research.



2

Literature review

2. Literature review

This chapter provides a detailed review of the existing literature related to this research. Firstly, the review explores the Charging Station Location Problem (CSLP) (see section 2.1) in a one-way car-sharing system, discussing various optimization methods, including traditional approaches, robust optimization, and stochastic programming. Furthermore, methods for traffic demand prediction are examined (see section 2.2), highlighting both classical techniques and deep learning approaches. Lastly, section 2.3 shows the research gaps in the existing literature and summarizes the main contributions of this thesis.

2.1. Charging Station Location Problem in One-Way Car-Sharing Systems

The CSLP in a one-way car-sharing system is a critical problem due to the increasing adoption of EVs and the need for accessible, efficient charging infrastructure. Various optimization methods have been applied to address this problem, including traditional methods, robust optimization, and stochastic programming. Among these, robust optimization and stochastic programming explicitly consider uncertainties such as fluctuating traffic demand. Robust optimization seeks solutions that perform well under various scenarios without assuming specific probability distributions, while stochastic programming incorporates probabilistic models of uncertainty. These approaches offer more resilient solutions compared to traditional methods that often assume deterministic conditions. In this section, we will classify and review relevant literature based on the optimization methods employed.

2.1.1. Traditional Methods

Traditional optimization methods are often the first approach applied to charging station location problems, using deterministic models and simpler mathematical formulations. A widely used method is Mixed-Integer Linear Programming (MILP), which provides exact solutions by formulating the problem through a set of linear equations and integer constraints. For example, [Deza et al. \[2018\]](#) applied MILP to optimize the placement of charging stations, enhancing vehicle flow balance and reducing relocation needs using a column generation technique to increase computational efficiency. Another powerful technique is the Continuum Approximation (CA) model, which is particularly effective in large-scale urban systems. [Li et al. \[2016\]](#) used a CA model to minimize costs related to station construction and vehicle relocation, solving NP-hard problems by decomposing large areas into manageable neighborhoods. An analytical approach was explored by [Bayram and Bayhan \[2020\]](#), who developed a theoretical framework using a traffic assignment technique with Stochastic User Equilibrium (SUE) to optimize the location of EV charging stations. Furthermore,

2. Literature review

game theory has also been utilized to optimize the location of charging stations by considering the strategic interactions between multiple stakeholders, such as station operators and EV users. [Meng and Kai \[2011\]](#) introduced a game-theoretic approach to enhance the decision-making process in station location planning. Their model considers the competitive environment, providing a more rational and scientifically sound framework for optimizing station locations. Multi-Objective Optimization methods are widely applied in this field. For instance, [Chen et al. \[2018\]](#) developed a Multi-Objective Particle Swarm Optimization (MOPSO) method that considers factors like land price, service distance, and installation capacity to minimize total costs while maximizing service capacity. Their approach effectively accounts for population density and land costs, making it a robust method for addressing the location-allocation problem for charging stations. Another traditional approach involves Genetic Algorithms (GA), which are frequently used for their ability to handle complex optimization problems. [Vazifeh et al. \[2015\]](#) applied a GA to optimize the deployment of EV charging stations in Boston. Their model significantly reduced drivers' excess driving distance to charging stations, energy overhead, and the number of stations required, demonstrating the practical effectiveness of genetic algorithms in real-world scenarios.

2.1.2. Robust Optimization (RO)

Robust optimization is used to handle the uncertainties in optimization problems by finding solutions that remain effective under various scenarios. This method is particularly effective in addressing the challenges posed by fluctuating demand, variability in energy supply, and the stochastic nature of traffic patterns, which traditional optimization methods often fail to handle adequately. Robust Optimization is designed to create feasible solutions under a wide range of uncertain scenarios. [Baron et al. \[2011\]](#) applied RO to a multi-period facility location problem, incorporating uncertainties in demand within bounded and symmetric multi-dimensional spaces. Their study demonstrated that different uncertainty models (box and ellipsoidal) lead to varying solution typologies, with RO providing more resilient solutions compared to deterministic approaches. In the context of EV charging stations, [Li et al. \[2022\]](#) developed a robust optimization model that integrates renewable energy sources and energy storage systems (ESS) into the planning of charging stations. The model accounts for uncertainties in both EV charging demand and renewable energy output. By using kernel density estimation, they addressed the issue of over-conservatism often associated with robust optimization, resulting in a more balanced and cost-effective infrastructure plan. A significant application of RO is seen in the work of [Wang et al. \[2019a\]](#), who proposed a hybrid model combining flow-refueling location modeling (FRLM) with robust optimization. This model effectively considers the uncertainty in charging demand and integrates queuing theory to optimize both the location and number of EV charging stations. The results indicated that the robust optimization approach, combined with queuing theory, significantly reduces the costs and enhances the efficiency of EV infrastructure deployment under uncertain demand conditions. Another approach was introduced by [Deb et al. \[2021\]](#), who developed a two-stage robust planning model that addresses the placement of charging stations while considering road traffic uncertainties. Their model utilizes a Bayesian network to model traffic randomness and applies a hybrid optimization algorithm to achieve an optimal balance between cost, voltage stability, and accessibility. This method ensures that the charging station network is resilient to fluctuations in traffic and demand.

2.1.3. Stochastic Programming

Stochastic programming addresses uncertainty in parameters such as demand and driving range in the CSLP by considering different possible scenarios, making the system more adaptable to real-world changes. Unlike robust optimization, which focuses on finding solutions that work well in the worst-case scenario, stochastic programming balances performance across all possible situations. In the following sections, we will first examine the uncertainty parameters, followed by the uncertainty modeling approaches, and finally, the stochastic programming approaches. This will help illustrate how stochastic programming effectively resolves uncertainty issues in the CSLP.

Uncertain parameters

The uncertainties addressed by stochastic programming primarily include the following aspects:

- **Demand Uncertainty:** This arises from fluctuations in EV charging demand across different time periods and locations. This variability in demand makes it challenging to predict how much charging infrastructure is required at specific locations. This is a common type of uncertainty considered in many types (see [Brandstätter et al., 2017], [Çalik and Fortz, 2019], [Hua et al., 2019], [Kim et al., 2021], [MirHassani et al., 2020]).
- **Driving Range uncertainty:** This refers to the variability in how far an EV can travel before getting a full charge. This type of uncertainty is influenced by various factors, such as battery conditions, road conditions, driving habits, and weather (see [Boujelben and Gicquel, 2019], [Wu and Sioshansi, 2017], [Jiao et al., 2017]).
- **Vehicle arrival time uncertainty:** This uncertainty concerns when EVs will arrive at charging stations, which affects station utilization and the planning of charging schedules. Variations in traffic, driving patterns, and user preferences lead to unpredictability in vehicle arrival times [Faridimehr et al., 2017].
- **State of charge uncertainty:** Variability in the battery levels of EVs arriving at charging stations, which affects the time and resources needed to recharge the vehicles [Faridimehr et al., 2017].
- **Charging behavior uncertainty:** Users' charging preferences and habits, such as when and where they choose to charge their vehicles. This behavioral uncertainty influences the utilization rates of charging stations and the overall demand patterns [Kim et al., 2021].

Uncertainty Modelling Approaches:

Uncertainty modeling plays a pivotal role in ensuring that decisions remain robust across various possible future scenarios. This process includes scenario generation and the assignment of probability distributions.

Scenario generation: This is one of the most prevalent approaches for modeling uncertainty in stochastic programming. In this method, multiple future scenarios are created based on different realizations of uncertain parameters. Each scenario represents a distinct possible future, and the goal is to optimize decisions that perform well across all scenarios (see

[Faridimehr et al., 2017], [Boujelben and Gicquel, 2019], [Brandstätter et al., 2017], [Çalik and Fortz, 2019], [Kim et al., 2021], [Fan, 2014]). This approach captures the variability in uncertain parameters in a structured way, making it easier to test solutions against a range of possible futures.

Probability Distributions: This approach is to model uncertainties using probability distributions. This method involves assigning probabilities to different outcomes for uncertain parameters. Decisions are then optimized by considering the expected values of these random variables. For example, Faridimehr et al. [2017] used probability distributions to model uncertainties in vehicle arrival times and state of charge. This approach provides a more direct and probabilistic understanding of the uncertainties, however, accurately estimating the probability distributions for certain uncertainties can be difficult, especially when historical data is limited.

Stochastic Decision Frameworks

A critical aspect of stochastic programming is the decision-making framework employed, which determines how uncertainty is incorporated into the model and how decisions are structured over time. From the perspective of decision-making stages, there are primarily two types of stochastic decision frameworks used in EV charging station planning: Two-Stage Stochastic Programming and Multi-Stage Stochastic Programming.

Two-Stage Stochastic Programming: This is a commonly used framework in the literature. Decisions are divided into two stages: in the first stage, strategic decisions such as the location of EV charging stations are made, while the second stage deals with operational adjustments after uncertainties are revealed. For example, Faridimehr et al. [2017] proposed a two-stage stochastic programming model to design an optimal network of charging stations, taking into account uncertainties in vehicle arrival times, state of charge, and drivers' walking ranges. The study utilized the Sample Average Approximation method, which converges asymptotically to an optimal solution, though the computational demands were significant. To overcome this, the authors also developed a heuristic to provide near-optimal solutions more efficiently. Similarly, Boujelben and Gicquel [2019] applied two-stage stochastic programming to optimize the placement of fast charging stations under driving range uncertainty. Their study tackled the stochastic flow refueling location problem (SFRLP), which posed considerable computational challenges, especially for large problem instances. To address this, they developed a Mixed-Integer Linear Programming (MILP) formulation along with a tabu search heuristic. The combination of these methods significantly reduced computation times while maintaining the quality of the solutions, making their approach viable for real-world applications. In another study, Kim et al. [2021] employed a two-stage stochastic mixed-integer programming (TSMIP) model to optimize the location of public charging stations. In the first stage, the locations of charging stations were determined, while the second stage involved the allocation of charging demand based on various scenarios reflecting uncertainties in users' charging behavior. Their model used dynamic decision trees to generate charging demand scenarios, offering a realistic representation of plug-in electric vehicle (PEV) users' behaviors. Brandstätter et al. [2017] also utilized a two-stage stochastic programming framework to optimize the placement of charging stations in Vienna. Their model incorporated demand uncertainty, and to handle the large-scale nature of the problem, the authors implemented heuristic methods, which enabled the model to solve problems of significant size without sacrificing solution quality. Additionally, Çalik and Fortz [2019] employed a two-stage framework, leveraging Benders decomposition to efficiently

solve demand-driven charging station location problems. Benders decomposition allowed them to break down the problem into smaller, more manageable subproblems, making their model particularly well-suited for large-scale applications. This approach also improved the computational feasibility of solving multiple demand scenarios in urban environments.

Multi-Stage Stochastic Programming: In multi-stage stochastic programming, decisions are made sequentially over multiple periods, allowing for ongoing adjustments as uncertainties evolve over time. This framework is particularly useful in long-term planning where uncertainties unfold gradually, and decisions can be adapted at different points in time. For example, Fan [2014] explored the application of multi-stage stochastic programming to optimize vehicle allocations in one-way car-sharing systems, accounting for fluctuating demand patterns. The study illustrated how the inclusion of demand uncertainty over multiple stages can significantly enhance decision-making processes, providing a more resilient system in the face of unpredictability. Similarly, Hua et al. [2019] proposed a multi-stage stochastic model for joint optimization of infrastructure planning and fleet management in EV sharing systems under uncertain demand. This approach highlighted the advantages of dynamic adjustments in fleet allocation and charging station placement, ultimately improving operational efficiency and service levels in highly volatile environments. Moreover, Kadri et al. [2020] introduced a multi-stage stochastic integer programming approach that incorporates a multi-period decision-making horizon, capturing the temporal dimension of uncertainties in recharging demand and infrastructure development. This method demonstrated how incorporating time-varying factors can lead to more robust and adaptive charging station networks.

2.1.4. Key Findings and Implications

Although traditional methods, robust optimization, and stochastic programming have been applied in addressing the CSLP in existing literature, each of these approaches has its limitations:

- Traditional methods often rely on deterministic models, which may not adequately capture the inherent uncertainties in real-world scenarios.
- Robust optimization, while effective in handling worst-case scenarios, can sometimes lead to overly conservative solutions that may not be optimal under typical conditions.
- Stochastic programming can address uncertainty, however, current stochastic programming models face challenges when dealing with demand uncertainty in the CSLP. It typically relies on assumed probability distributions or predefined scenarios, which may fail to fully capture the complex and dynamic nature of traffic demand. As the number of scenarios increases to better represent uncertainty, the computational complexity rises significantly, potentially limiting its application in large-scale real-time problems.

This research focuses on demand uncertainty, which is one of several types of uncertainties. The rapid advancement of data-driven approaches, especially in deep learning, now allows us to use historical traffic data to produce more accurate demand forecasts. This significantly enhances the decision-making capabilities of stochastic programming models. Therefore, a detailed exploration of traffic demand forecasting methods, particularly deep learning models, will offer a strong data foundation for optimizing charging station locations and improving overall system efficiency.

2.2. Methods for Traffic Demand Prediction

Traffic demand forecasting is critical in transportation research, particularly in the CSLP for one-way car-sharing systems, as it directly impacts the accuracy of infrastructure planning and optimization outcomes. In the following section, we will review several methods of traffic demand forecasting, with a specific focus on how deep learning models play a pivotal role in solving the CSLP.

2.2.1. Classical Methods

In this domain, traditional statistical and machine-learning models play a significant role. Statistical models are widely used for short-term traffic prediction due to their ability to treat forecasting as a regression problem, offering clear computational structures and strong theoretical interpretability. [Chrobok et al. \[2004\]](#) analyzed daily and seasonal traffic patterns using two years of data from inductive loops in Duisburg, Germany. They then developed an automatic matching process to classify new traffic data. Subsequently, they used constant and linear models to predict short-term traffic. The research found that combining short-term and long-term forecasting methods improves traffic prediction accuracy. In time series analysis, Autoregressive Integrated Moving Average (ARIMA) models are widely used due to their ability to handle traffic demand data effectively. [\[Yu and Zhang, 2004\]](#) proposed a switching ARIMA model for short-term traffic flow forecasting, effectively capturing the dynamic patterns in traffic data. It demonstrates superior performance over traditional models, especially in handling transitions between different traffic states. The model's application to real-world data from Beijing confirms its effectiveness and potential for broader use in ITS. [Alghamdi et al. \[2019\]](#) addresses the growing problem of traffic congestion by leveraging ARIMA. They focused on non-Gaussian traffic data and provided a systematic approach to preprocessing and modeling using real traffic data from California. The results demonstrate that ARIMA (2,1,3) closely matches actual traffic patterns, offering a reliable method for managing traffic congestion. Seasonal ARIMA (SARIMA) is an extension of the ARIMA that specifically accounts for seasonality in time series data. It incorporates seasonal components in addition to the non-seasonal components used in a standard ARIMA model. The study by [Kumar and Vanajakshi \[2015\]](#) explored the use of SARIMA models for predicting short-term traffic flow with minimal data. Using data from only three consecutive days on a 3-lane arterial roadway in Chennai, India, the SARIMA model achieved a mean absolute percentage error of 4-10%, which is acceptable for Intelligent Transportation Systems applications.

Traditional machine learning approaches, which can handle more complex data, are broadly divided into three categories: feature-based models, Gaussian process models, and state space models [\[Shaygan et al., 2022\]](#). The comparison of these methods is shown in Table 2.1. Feature-based methods are used to train regression models, including linear regression for various traffic prediction problems using human-engineered traffic features (see [\[Guan et al., 2018; Li et al., 2019\]](#)). Gaussian process models are used to predict traffic demand by modeling the distribution over possible functions that fit the observed data. These models leverage the properties of Gaussian distributions to provide a flexible and powerful method for making predictions based on spatial and temporal data patterns (see [\[Diao et al., 2019; Salinas et al., 2019\]](#)). State space models can model the system's dynamics with hidden states that evolve over time according to a set of linear or nonlinear equations. These models capture temporal dependencies and uncertainties. For example, [Zhou and Mahmassani](#)

2. Literature review

[2007] introduced a structural state space model for dynamic OD demand estimation. It integrates regular demand patterns, structural deviations, and random fluctuations using a polynomial trend filter and Kalman filtering for real-time updates. The model is designed to improve real-time dynamic traffic assignment (DTA) systems by continuously utilizing real-time traffic data. Similarly, Pan et al. [2013] enhances the Stochastic Cell Transmission Model by incorporating spatial-temporal correlations to improve short-term traffic state prediction. By integrating a multivariate normal distribution-based predictor and calibrating the covariance structure from spatial correlation analysis, the model captures the dynamic dependencies in traffic flow. Predictions are conducted in a rolling horizon manner, allowing real-time adjustments and enhancing accuracy. Furthermore, Duan et al. [2019] proposed a unified spatio-temporal model for short-term traffic flow prediction, capturing the time-varying spatio-temporal correlations between traffic at different measurement points. The model is physically intuitive, considering road network topology, time-varying speed, and trip distribution, distinguishing it from previous black-box approaches. By integrating physical factors affecting spatio-temporal correlation into adjustable parameters, the model reduces computational complexity and adapts easily to changing road and traffic conditions. Experimental results using two real traffic datasets demonstrate superior accuracy compared to traditional STARIMA and neural network approaches, with reduced computational complexity.

Table 2.1.: Comparison of three traditional machine learning methods

Criteria	Feature-Based Methods	Gaussian Process Models	State Space Models
Simplicity and Understandability	Simple and easy to understand	Less intuitive to interpret	Requires strong assumptions about the underlying process
Flexibility	Flexible with a wide range of algorithms (e.g., SVM, Decision Trees)	Non-parametric and flexible, can model complex functions	Naturally models temporal and sequential data
Computational Efficiency	Often computationally efficient, especially with large datasets	Computationally expensive, especially with large datasets	Computationally intensive, especially for large state spaces
Uncertainty Estimates	Generally does not provide uncertainty estimates	Provides uncertainty estimates, valuable for understanding prediction confidence	Provides a probabilistic framework for inference and prediction
Interpretability	Easy to interpret results	Less intuitive and harder to interpret	Interpretation can be challenging due to model complexity

2.2.2. Deep learning Methods

Deep learning has been regarded as a powerful approach to predict traffic demand in recent years. They effectively address the limitations of traditional methods, which struggle with capturing complex, nonlinear patterns in traffic data and require strong assumptions about




Figure 2.1.: Multilayer Perceptron (MLP)

data behavior. Initially, deep learning approaches were seen as impractical due to their high demands for data, computation, and storage compared to classical methods. However, with the advent of vast data availability and increased computational power, deep learning has shown significant potential in traditional methods in traffic prediction [Shaygan et al., 2022]. Next, we will explore traffic demand forecasting by examining four common deep-learning methods.

Multilayer Perceptron(MLP)

Multilayer Perceptron (MLP) is one of simplest forms of DNNs, they consist of input, hidden, and output layers, where each neuron in one layer is connected to every neuron in the next later (see Figure 2.1). MLPs can capture nonlinear relationships in traffic data and are often used as baseline models in traffic prediction studies [Gao et al., 2020]. Qin et al. [2023] proposed a novel spatiotemporal hierarchical MLP network (STHMLP) for traffic forecasting which addresses the limitations of traditional traffic forecasting methods by capturing hierarchical temporal characteristics and macro spatial dependencies. The model employs a decomposition architecture with fine and coarse modules to extract detailed spatio-temporal information from road and region networks. Extensive experiments on four real-world traffic datasets demonstrate that STHMLP significantly outperforms existing state-of-the-art methods in terms of accuracy and efficiency. Dimara et al. [2021] presented a method for predicting traffic volumes at toll plazas using a multi-layer perceptron (MLP) neural network. By integrating minimal yet effective data such as vehicle counts, weather conditions, and temporal features, the model achieves high accuracy with a mean absolute percentage error (MAPE) of approximately 8.85%. The method is cost-effective and non-intrusive, suitable for real-time traffic management, thus aiding in reducing congestion and enhancing road network efficiency. Specifically for the charging station location problem in one-way car-sharing systems, deep learning methods can generate more accurate traffic demand forecasts, thereby providing high-quality scenario data for stochastic programming models. This, in turn, enhances the reliability and robustness of the optimization process.

Convolutional Neural Networks (CNNs)

Convolutional Neural Networks have been increasingly applied to traffic prediction tasks due to its ability to capture spatial and temporal features from traffic data. Since it is challenging to represent the traffic network with 2D matrices, several researchers have tried converting the traffic network structure at different times into image. These images are then divided into standard grids, with each grid representing a specific region [Yin et al., 2021]. Ma et al. [2017] introduced a novel approach that transforms spatiotemporal traffic dynamics into images, enabling CNNs to predict traffic speed across large-scale networks. This method involves creating a two-dimensional time-space matrix to represent traffic flow, which is then processed by a CNN for feature extraction and prediction. The study showed that this method outperforms traditional algorithms, such as ordinary least squares and k-nearest neighbors, as well as other deep learning models like stacked auto-encoders and long short-term memory networks (LSTMs), with a substantial improvement in accuracy. Furthermore, Duan et al. [2018] presents a novel approach to enhancing urban traffic flow prediction through use of a deep hybrid neural network. They combine Convolutional Neural Networks (CNNs) and Long Short-Term Memory networks (LSTMs) to effectively capture the spatial and temporal features of traffic data. The model is further optimized using a greedy algorithm, which improves prediction accuracy and reduces computation time. The approach was validated using real-world taxi GPS data from Xi'an, demonstrating superior performance compared to traditional prediction methods. To address the spatial and temporal complexities of traffic demand, Du et al. [2020] proposed the Dynamic Transition Convolutional Neural Network (DTCNN). This model was specifically designed to capture the spatial distributions and the evolving dynamics of traffic demand by constructing a transition network based on historical data. The DTCNN was tested on New York City taxi and bike-sharing data, where it demonstrated its effectiveness in producing precise traffic demand predictions. Recent advancements have further extended the application of CNNs in traffic prediction by integrating graph-based approaches. For example, Chen et al. [2020] introduced a dynamic spatio-temporal graph-based CNN (DST-GCNN) that learns expressive features to represent spatio-temporal structures and predict future traffic flows. The DST-GCNN model utilizes a novel graph-based spatiotemporal convolutional layer to capture dynamic relationships in traffic data, showing competitive performance against other state-of-the-art methods. This approach highlights the potential of CNNs to not only handle static spatial information but also adapt to the dynamic nature of traffic networks over time.

Recurrent Neural Networks (RNN)

Recurrent Neural Networks (RNN), particularly Long Short-Term Memory (LSTM) networks and Gated Recurrent Units (GRUs), have shown substantial promise in traffic prediction tasks due to their ability to learn and model temporal dependencies from sequential data. These models are particularly well-suited for forecasting traffic flow, speed, and demand, where the sequence of past data points plays a critical role in predicting future trends. For instance, Vinayakumar et al. [2017] applied various RNN architectures, including LSTM and GRU, to network traffic prediction, demonstrating that these models outperform classical methods in capturing long-term dependencies and providing accurate traffic predictions. Similarly, Ramakrishnan and Soni [2018] explored the use of RNNs for network traffic prediction, comparing the performance of standard RNNs, LSTMs, and GRUs across different

2. Literature review

traffic prediction tasks. Their study found that RNN-based models significantly outperformed traditional statistical methods, particularly in tasks like volume prediction, packet protocol prediction, and packet distribution prediction. Further advancements have been made by integrating RNNs with other deep-learning models to enhance prediction accuracy. For example, [Zhu et al. \[2020\]](#) proposed a novel traffic flow prediction method combining RNNs with Graph Convolutional Networks (GCNs) and a belief rule base, which effectively captures both spatial and temporal correlations in traffic data. Moreover, the work by [Tokuyama et al. \[2018\]](#) highlighted the importance of incorporating additional contextual information, such as timestamps and days of the week, into RNN models. Their study showed that using such attribute information alongside traditional time-series data could significantly improve the prediction accuracy of RNNs for network traffic prediction.

Graph Convolution Network (GCN)

Graph Convolutional Networks (GCNs) have emerged as a powerful tool for traffic prediction due to their ability to model the spatial dependencies in traffic networks effectively. These networks represent traffic systems as graphs where nodes correspond to specific locations (e.g., road segments, stations) and edges represent the spatial relationships between these locations.

One prominent approach is the Temporal Graph Convolutional Network (T-GCN), which integrates GCN with Gated Recurrent Units (GRUs) to capture both spatial and temporal dependencies in traffic data. [Zhao et al. \[2018\]](#) demonstrated that the T-GCN model outperforms traditional methods by effectively learning the complex topological structures of urban road networks and the dynamic changes in traffic flow. Further advancements in GCN-based traffic prediction include the development of models that dynamically adjust the graph structure to better reflect the changing spatial relationships in traffic data. For instance, [Guo et al. \[2020\]](#) introduced a Dynamic Graph Convolution Network (DGCN) that constructs dynamic road network graphs by learning latent spatial-temporal features. This approach enhances the model's ability to adapt to the non-stationary nature of traffic data, leading to improved prediction accuracy. Another innovative approach is the Coupled Layer-wise Graph Convolution model proposed by [Ye et al. \[2021\]](#), which uses self-learned adjacency matrices to capture multi-level spatial dependencies within the network. This model has shown superiority over traditional methods by allowing different layers to learn distinct spatial relationships, further improving the accuracy of transportation demand predictions.

Table 2.2.: A Review of Deep Learning Methods in Traffic Demand Prediction Research

Author	Dataset	Spatial modeling type	Temporal modeling type	Evaluation Criteria
Qin et al. [2023]	California (PeMS)	Multilayer Perceptron	-	MAE, RMSE, MAPE
Dimara et al. [2021]	Analipsi, Greece(toll plaza data)		Multilayer Perceptron	MAE, MAPE
Ma et al. [2017]	Beijing (GPS data)	CNN	CNN	MAE, RMSE, R^2
Duan et al. [2018]	Xi'an, China (GPS data)	CNN	RNN(LSTM)	MSE, RMSE
Du et al. [2020]	NYC dataset	CNN	-	MAE, RMSE
Chen et al. [2020]	Bike-NYC, Taxi-NYC	CNN (DST-GCNN)	CNN (DST-GCNN)	MAE, RMSE, MAPE
Venayakumar et al. [2017]	GÉANT network data	-	LSTM, RNN, GRU	RMSE, MSE
Ramakrishnan and Soni [2018]	GÉANT, Abilene (network data)	-	RNN, LSTM, GRU	RMSE, MSE
Zhu et al. [2020]	Shenzhen Taxi Dataset	GCN	RNN	RMSE, MSE
Tokuyama et al. [2018]	WIDE Project Dataset	-	RNN	RMSE, MSE
Zhao et al. [2018]	Shenzhen-taxi dataset and Los-loop dataset	GCN	GRU	RMSE, MSE
Guo et al. [2020]	PeMSD4, PeMSD8, and PHILADELPHIA	GCN(DGCN)	-	RMSE, MSE
Ye et al. [2021]	NYC City Bike, NYC Taxi	GCN	GRU	RMSE, MSE

To summarize the research outcomes of existing traffic demand forecasting models, Table 2.2 outlines the application of various studies across different datasets, as well as their spatial and temporal modeling methods and evaluation metrics. Building on this, this research will adopt the LSTM-MLP-MDN architecture to address this issue. This architecture not only captures complex spatiotemporal features but also generates more accurate demand probability distributions, thereby improving the optimization performance of stochastic programming models under uncertainty. Compared to the methods reviewed in Table 2.2, LSTM is effective in handling time series data, while MLP can model nonlinear features, particularly the contextual information of traffic data. MDN further enhances the modeling of the probability distribution of traffic demand, enabling more precise predictions of demand fluctuations. In addition, this research will use the New York City Taxi dataset as the research data, employing MAE, MAPE, and R^2 as the primary evaluation metrics. These indicators will help assess the performance of the proposed method in predicting traffic demand across different time periods and spatial locations and will facilitate direct comparisons with existing models to validate the effectiveness and robustness of the approach in real-world applications.

2.3. Research Gap and Main Contributions

Many existing studies have made large advances in the modeling and optimization of charging station location problems for electric car-sharing services, primarily utilizing deterministic models that do not incorporate uncertainty. However, in this problem, demand is highly dynamic, and such models may oversimplify the complexities of real-world scenarios by not adequately addressing demand uncertainties inherent in these systems.

With the consideration of demand uncertainty, some researchers have explored stochastic programming models to address this problem [Brandstätter et al., 2017; Çalık and Fortz, 2019], these approaches, generate multiple scenarios to approximate real-life demand patterns. It is worth noting that the probability distribution of traffic demand has a great impact on the quality of this problem solution. There are basically two ways of dealing with it in stochastic programming. One way is to utilize the expectation of random variables to replace the probability distribution, in this case, the stochastic programming model practically becomes a deterministic optimization model. The other way is to use parametric approaches, where the probability distribution is assumed. Brandstätter et al. [2017] predefined seven scenarios based on the days of the week with determined probabilities. Çalık and Fortz [2019] utilized probability distribution functions fitted from historical data via a nonlinear least square optimization method to generate scenarios, each with a specific number of trips, then compute the probability of scenarios based on the probability of individual trip in each scenario. However, their approaches still exhibit limitations in adaptability and robustness in infrastructure planning, they may not fully capture the complexity of demand patterns, which can fluctuate significantly across different times. Consequently, there remains a critical need to explore methodologies that can deliver a more accurate and dynamic approach to demand forecasting, thereby enhancing the reliability and effectiveness of charging station location optimization for electric car-sharing services.

Motivated by this gap in the literature, we propose a novel deep learning-based stochastic programming framework (LMSP Framework) to address the CSLP. Our model integrates Long Short-Term Memory (LSTM) networks, Multilayer Perceptron (MLP), and Mixture Density Networks (MDN). The LSTM component is employed to capture and forecast the

2. Literature review




Figure 2.2.: Research gap

temporal dynamics of traffic demand, while the MLP serves to enhance the predictive power by capturing complex, non-linear relationships within the contextual data. The MDN is then utilized to derive the probability distribution of the predicted traffic demand, enabling the generation of various demand scenarios. These scenarios are subsequently used within a stochastic programming model, which strategically optimizes the locations and capacities of charging stations under demand uncertainty, ultimately aiming to maximize system profitability while ensuring robust and reliable infrastructure planning.


To the best of our knowledge, this is the first research to utilize a deep learning-based stochastic programming framework to address CSLP in a one-way electric car-sharing system while explicitly considering traffic demand uncertainty. By employing a two-stage stochastic programming approach, we effectively handle demand uncertainty. The application of the deep learning model (LSTM-MLP-MDN) is used to enhance the two-stage stochastic programming model because it allows for more accurate forecasting of traffic demand probability distribution. This approach represents a significant advancement in this field, offering a more robust and reliable method.

The main contributions of this research are as follows:

1. **Novel LMSP Framework:** This research presents a novel framework that integrates CSLP with stochastic programming and data-driven methods. This approach directly addresses the identified research gaps (see Figure 2.3), offering a structured solution to the challenges of optimizing charging infrastructure under uncertainty.
2. **Data-Driven Enhancement of Stochastic Programming:** This research improves the existing two-stage stochastic programming model by integrating a deep learning model specifically designed to predict traffic demand probability distributions. By using historical traffic data, the deep-learning model captures complex demand patterns and provides more accurate inputs for the stochastic programming model.
3. **LSTM-MLP-MDN model:** A new deep learning architecture, combining Long Short-Term Memory (LSTM), Multilayer Perceptron (MLP), and Mixture Density Networks (MDN), has been developed. This hybrid model enables the prediction of traffic demand distributions, addressing uncertainties in the CSLP. By learning from historical

2. *Literature review*

data, the model generates realistic demand scenarios, improving the reliability of infrastructure planning.

A low-angle, blue-tinted photograph of the Statue of Liberty. The statue is the central focus, holding a torch in her right hand and a tablet in her left. The background is a clear blue sky with some light clouds. The overall tone is monochromatic and professional.

3

Problem Description and Formulation

3. Problem Description and Formulation

This chapter presents a detailed description of the problem this research aims to address, followed by the formulation of our proposed solution. Section 3.1 describes the problem that this research seeks to address. Section 3.2 introduces the two-stage stochastic programming model for the charging station location problem. This model is primarily based on [Brandstätter et al., 2017; Çalik and Fortz, 2019], with some refinements applied to better align it with real-world conditions.

3.1. Problem Description

The rapid growth of electric vehicles (EVs) and car-sharing services has introduced significant challenges in urban transportation planning, particularly in optimizing the location of charging stations for one-way electric car-sharing systems under demand uncertainty. This issue is crucial for promoting sustainable urban mobility, but it remains a significant challenge due to the highly dynamic nature of traffic demand and the inherent complexities of urban environments.

In one-way electric car-sharing systems, users can pick up a vehicle at one location and return it to another, offering flexibility but complicating resource management. The core problem addressed in this research is how to strategically locate charging stations and allocate electric vehicles to maximize system profitability while maintaining high service levels. The key challenge is the demand uncertainty, influenced by factors such as time of day, weather conditions, and unforeseen events, which complicates the prediction of when and where vehicles will be needed.

To address this, the research formulates the problem as a two-stage stochastic programming model. The first stage involves strategic decisions, such as where to place charging stations and how many vehicles to allocate to each station. The second stage focuses on operational decisions, such as whether to accept a user's trip request based on vehicle availability and predicted demand, and whether a specific accepted trip is assigned to a particular car. Incorporating a deep learning model to forecast traffic demand, this research aims to enhance decision-making under demand uncertainty. The model combines the strengths of Long Short-Term Memory (LSTM) networks, Multilayer Perceptron (MLP), and Mixture Density Networks (MDN) to predict the probability distribution of future traffic demands. The predicted demand distribution serves as input for the stochastic programming model, significantly improving the robustness and reliability of traffic management strategies.

3.2. Two-Stage Stochastic Programming Model

Firstly, it is necessary to identify the potential locations for charging stations in the city's popular urban areas, denoted as I . Subsequently, we aim to select optimal charging stations in I and assign an initial number of vehicles to the selected stations. In this system, users from different regions request trips, but these requested trips need to first satisfy certain conditions to be accepted and then receive permission to use the electric vehicle car-sharing service.

3.2.1. Assumptions

To accurately model this system, it is crucial to establish certain assumptions before formulating the two-stage stochastic programming model. These assumptions are as follows:

- **No Operational Activities:** The car-sharing company will not consider the operational activities of the service staff, such as car relocation fees and charging station relocation fees. This assumption ensures the decision of the first stage is fixed.
- **Uniform Vehicle Type:** All electric cars located in the station conform to a uniform type, indicating identical automotive performance characteristics and battery capacity.
- **Battery Consumption:** The battery consumption of electric cars is directly proportional to their travel time, meaning that all vehicles have the same maximum travel time. This helps to reduce the impact of uncertainty in battery levels on the analysis.
- **Charging Station Capacity:** Each potential charging station location has the same maximum capacity, with all selected stations equipped with the maximum number of charging slots (also referred to as charging parking spaces). This can prepare for peak demand conditions in high-traffic areas.
- **OD Demand:** All traffic demand starts and finishes at potential charging stations. This simplifies the process for users to reach the nearest charging station.
- **User Behavior:** Users will choose the shortest travel time route.
- **Charging Requirements:** Initially, the electric vehicles at the station are fully charged to 100%. Each vehicle must be fully charged before use. Once an electric vehicle is utilized and returned to the charging station, it starts recharging immediately. The charging time for all vehicles is the same, and upon completion, the battery capacity will again reach 100%.
- **No Queue for Parking:** The need to queue for returning vehicles is not considered. This can eliminate potential delays or congestion at charging stations.

3.2.2. Notations and Definitions

Next, we will introduce some important notations and definitions in the optimization model, which can provide clarity and precision in our model formulation. Here, upper case letters represent sets and lower case ones represent indices. For example, I is a set of potential stations, and i is the index of an element of I .

3. Problem Description and Formulation

- **Sets**

- $I = \{1, 2, \dots, |I|\}$: The set of potential locations for charging stations, where i indexes the stations.
- $T = \{0, 1, 2, \dots, T_{max}\}$ is the set of time periods (of identical length) in the planning period, where t denotes indices of the set.
- $S = \{1, 2, \dots, |S|\}$ is the set of demand scenarios, where s denotes indices of the set.
- $H = \{1, 2, \dots, |H|\}$ is the set of purchased cars, where h represents indices of the cars.

- **Model Parameters and Coefficients**

- p_s is the probability of scenario $s \in S$.
- i is the income of each trip.
- f_i is the fixed cost of the station if built at location i .
- c is the fixed purchasing cost of each car.
- W is the budget limit.
- N is the highest number of vehicles that can be purchased.
- t_k^s is the shortest travel time required for the k -th trip, in scenario s .
- Φ is the maximum travel time of a fully charged car.
- C is the capacity of station i , $i \in I$.
- τ is the required charging time after a car is returned to the charging station.

Decision Variables

- **First Stage Decision Variables (Strategic Layer)**

- $y_i = 1$ if a charging station is built at location i , $\forall i \in I$, 0 otherwise.
- L_i is the initial number of electric vehicles at station i , $\forall i \in I$.

- **Second Stage Decision Variables (Operational Layer)**

- $x_k = 1$ if trip k is accepted, $\forall k \in K^s$, $\forall s \in S$, 0 otherwise.
- $x_k^h = 1$ if an accepted trip k of scenario s is assigned to car h , $\forall h \in H$, $\forall s \in S$, $\forall k \in K^s$, 0 otherwise.
- $f_a^h = 1$ if car h is traveling along a particular arc a , $\forall h \in H$, $\forall a \in A^s$, 0 otherwise.

In the planning period, the set S includes various scenarios of total traffic demand, reflecting the traffic demand uncertainty. The probability of each corresponding scenario is denoted by p_s . Each demand scenario $s \in S$ includes a set of requested trips K^s , where each requested trip $k \in K^s$ consists of four elements: an origin point o_k , a starting time s_k , a destination point d_k , and an ending time e_k . Note that $o_k \in I$ and $d_k \in I$, and the travel time of trip k is computed as $e_k - s_k$.

In each demand scenario $s \in S$, a requested trip $k \in K^s$ can only be accepted if a purchased car h is assigned to it. For each purchased car h , $h \in H$, let $K_h^s = (k_{h_1}^s, k_{h_2}^s, \dots, k_{h_l}^s)$ be the sequence of trips carried out with car h in time order, where l represents the total number of trips performed by a car.

3. Problem Description and Formulation

The decision variables in this model are divided into two stages. The first stage is the strategic layer, where the optimal locations for the charging stations y_i and the initial number of electric vehicles assigned to each station L_i are determined. This stage involves making long-term decisions that significantly affect the overall efficiency and profitability of the car-sharing system. The second stage is the operational layer, which focuses on the short-term operations of the system. Here, there are three decision variables. x_k evaluates whether a user's requested trip is accepted based on the real-time availability of vehicles and the locations of built charging stations. x_k^h ensures that each trip is matched with an appropriate vehicle. Flow variables f_a^h track the movement of each car, ensuring that the distribution of cars aligns with the system's operational requirements.

Time-Expanded Location Graphs (TLG)

This model introduces a time-expanded location graph $G^s = (V^s, A^s)$ for each scenario s , which enables the monitoring of the position of each car at each time point (see Figure 3.1a as an example). Here, nodes represent the possible locations (stations) of vehicles at different time periods, and arcs represent the possible transitions (movements or status changes) of vehicles between these locations over time. The graph in this research is similar to the one used by Brandstätter et al. [2017], but there are a few modifications, such as in the initial allocation. Through practical implementation, it was found that replacing a large number of binary decision variables z_h , which indicate whether or not car h is purchased, with a smaller number of integer variables L_i can significantly reduce runtime and reduce memory issues in complex integer programming problems. Additionally, these adjustments were made to clearly show the number of electric vehicles allocated to each charging station, making it easier to understand.

Nodes

In the graph, the node set V^s consists of a sink node s^s , and each node i_t for each station $i \in I$ and each time point $t \in \{0, 1, 2, \dots, T_{max}\}$. Here, note that i is used to represent an index of a location in models that do not necessarily consider the time dimension explicitly. On the other hand, i_t is employed to explicitly denote a temporal aspect of a node, indicating the state of node i at time t .

Arcs

A^s is the set of arcs, which include waiting arcs, travel arcs, and final collection arcs:

- **Waiting Arcs** A_W^s : Waiting arcs are arcs that connect a station i at one time period to the same station at the subsequent time period $t + 1$. They represent the scenario where a vehicle remains at the same station, either parked or charging. Formally, these waiting arcs are defined as: $A_W^s = \{(i_t, i_{t+1}) \mid i \in I, t \in \{0, 1, 2, \dots, T_{max} - 1\}\}$
- **Travel Arcs** A_T^s : Travel arcs are arcs that represent the movement of a vehicle from station i to another station j over a specific time period. They can capture the actual movement of vehicles between different stations as they fulfill trip requests. These arcs corresponding to trips requested in the scenario and are formally defined as $A_T^s(k) = \{(i_{s_k}, j_{e_k}) \mid i \in I, j \in I, s_k \in T, e_k \in T, s_k < e_k\}$ for a trip $k \in K^s$, and travel arcs is $A_T^s = \bigcup_{k \in K^s} A_T^s(k)$.

3. Problem Description and Formulation

- **Final Collection Arcs** A_C^s : Final collection arcs connect each station i at the last time period T_{max} to a sink node s^s . These arcs ensure that the time-expanded graph is complete by providing a way to 'collect' all flows from the last time period. Furthermore, they ensure the model can respect the capacity of the system at the end of the planning horizon. Formally, these arcs are defined as $A_C^s = \{(i_{T_{max}}, s^s) \mid i \in I\}$.

In TLG, parallel travel arcs may exist because different trips might start and end at the same points within the same time period but be realized by different cars. Similarly, if cars are parked at the same station at the same time, their waiting arcs and final collection arcs may also be identical. But TLG does not contain parallel waiting or collection arcs, since each arc will be linked to each available car.

We further define:

- $\delta^-(u) = \{(v, u) \in A^s \mid \text{arc}(v, u) \text{ exists from node } v \text{ to node } u\}$
- $\delta^+(u) = \{(u, v) \in A^s \mid \text{arc}(u, v) \text{ exists from node } u \text{ to node } v\}$

$\delta^-(u)$ represent the set of incoming arcs of a node u . It can capture all possible movements or flows that lead into node u , whether they are travel arcs (representing vehicles moving into a station) or waiting arcs (representing vehicles that are waiting or charging). $\delta^+(u)$ represent the set of outgoing arcs of a node u . It can capture all possible movements or flows that originate from node u , which could be vehicles departing from a station to another station or vehicles moving to a different state within the same station (e.g. from waiting to traveling).

Furthermore, the notation $f^h[A'] = \sum_{a \in A'} f_a^h$ indicates the subset of flow variables f_a^h for a subset of arcs $A' \subset A^s$. The variable f_a^h represents the flow of car h on the waiting arc a' . Specifically, this variable indicates whether car h is located on the arc a' at a certain time period after completing a trip. The arc a' corresponds to the time period during which the vehicle is waiting or charging at a station before it is ready for the next trip.

Illustrative Example


Next, we explain the time-expanded location graph through diagrams with a simple example. As shown in 3.1a, in scenario S' , there are three potential charging station locations, all of which are built, denoted as $i \in \{1, 2, 3\}$. The planning period is $T = \{0, 1, T_{max}\}$. At $t = 0$, cars are initially allocated to each built charging station (see Figure 3.1b). All arcs reflect the activities of cars at different time periods and different points, and the sink node is used to aggregate and simplify constraints related to the flow. For example, we assume each station initially has one car h , where $h \in \{1, 2, 3\}$, and the car states $K_h^s = (k_{h_1}^{S'}, k_{h_2}^{S'})$ reflect the activities of each car. In Figure 3.1b, arcs of different colors represent the activities of vehicles. The green car travels from i_1 to i_2 between $t = 0$ and $t = 1$, and then parks at i_2 for charging from $t = 1$ to T_{max} . The yellow car parks from $t = 0$ to $t = 1$, then travels from i_2 to T_{max} . However, the red car remains parked throughout the planning period.

Then Table 3.1 summarizes the arcs appear in this example:


3. Problem Description and Formulation

Table 3.1.: Description of different types of arcs

Category	Arc Type	Arcs
Waiting arcs	(i_t, i_{t+1})	$(1_0, 1_1), (1_1, 1_{T_{max}}), (2_0, 2_1), (2_1, 2_{T_{max}}), (3_0, 3_1), (3_1, 3_{T_{max}})$
Travel arcs	(i_t, j_{t+1})	$(1_0, 2_1), (2_1, 1_{T_{max}})$
Final collection arcs	$(i_{T_{max}}, \text{sink})$	$(1_{T_{max}}, \text{sink}), (2_{T_{max}}, \text{sink}), (3_{T_{max}}, \text{sink})$



(a) An example of Time-expanded Location Graph



(b) Detailed Traffic Flow Graph

Figure 3.1.: Time-expanded Location Graph

3.2.3. Time-Dependent Integer Linear Program

The goal of this model is to find the optimal location for charging stations and determine the initial allocation of electric cars to these stations in a one-way car-sharing system under demand uncertainty. The objective function is established to maximize the profit of the car-sharing company by calculating the total income from services minus all associated costs. Now Time-Dependent Integer Linear Program is written as follows:

3. Problem Description and Formulation

$$\max \sum_{s \in S} p_s \sum_{k \in K^s} i x_k - \sum_{i \in I} f_i y_i - c \sum_{i \in I} L_i \quad (3.1)$$

$$\text{s.t.} \quad \sum_{i \in I} f_i y_i + c \sum_{i \in I} L_i \leq W \quad (3.2)$$

$$\sum_{i \in I} L_i \leq N \quad (3.3)$$

$$0 < t_k^s x_k \leq \Phi \quad \forall s \in S, \quad \forall k \in K^s \quad (3.4)$$

$$\sum_{h=1}^H \sum_{a \in \delta^+(i_0)} f_a^h = L_i y_i \quad \forall s \in S, \quad \forall i_0 \in V_0^s \quad (3.5)$$

$$\sum_{h=1}^H x_k^h = x_k, \quad \forall s \in S, \quad \forall k \in K^s \quad (3.6)$$

$$\sum_{h \in H} \sum_{k \in K^s: o_k = i, s_k = 0} x_k^h \leq L_i y_i, \quad \forall i \in I, \quad \forall s \in S \quad (3.7)$$

$$\sum_{h=1}^H \sum_{a \in \delta^+(i_t) \cap (A_W^s \cup A_C^s)} f_a^h \leq C y_i \quad \forall s \in S, \quad \forall i_t \in V^s \setminus \{s^s\} \quad (3.8)$$

$$f^h[\delta^-(i_t)] \leq y_i \quad \forall h \in \{1, 2, \dots, H\}, \quad \forall i_t \in V^s \setminus \{s^s\} \quad (3.9)$$

$$f^h[\delta^-(i_t)] = f^h[\delta^+(i_t)] \quad \forall h \in \{1, 2, \dots, H\}, \quad \forall i_t \in V^s \setminus \{s^s\} \quad (3.10)$$

$$\sum_{a \in A_T^s(k)} f_a^h = x_k^h \quad \forall h \in \{1, 2, \dots, H\}, \quad \forall s \in S, \quad \forall k \in K^s \quad (3.11)$$

$$f_a^h \leq f_{a'}^h \quad (3.12)$$

$$\forall h \in \{1, 2, \dots, H\}, \quad \forall s \in S, \quad \forall k \in K^s,$$

$$\forall a = (i_{s_k}, j_{e_k}) \in A_T^s(k),$$

$$\forall a' = (j_{e_k}, j_{e_k + \tau}) \in A_W^s$$

The objective function 3.1 maximizes the expected second-stage profit contribution from the accepted trips. The first term in the objective function represents the expected income from the car-sharing service. The second term accounts for the total fixed cost of all charging

3. Problem Description and Formulation

stations, while the third term reflects the purchasing cost of all vehicles.

Constraints 3.2 require that the combined total construction cost of the charging stations and the purchase cost of electric vehicles should not exceed the budget limit W . Constraints 3.3 ensure that the initial number of allocated vehicles does not surpass the total vehicle count. Constraints 3.4 ensure that trip k does not exceed the maximum travel time of the car Φ due to battery limitations. Constraints 3.5 impose restrictions on nodes i_0 , which represents node i of TLG in the initial state ($t = 0$), and V_0^s represents the nodes at the initial time $t = 0$ within the scenario s . Specifically, the sum of all arcs originating from built charging station i , including both waiting arcs and traveling arcs, must equal the initial number of cars at charging station i . These constraints can guarantee that each purchased car is first allocated to its corresponding built charging station. Constraints 3.6 ensure that exactly one car is assigned to each accepted trip. Constraints 3.7 ensure that the total number of trips assigned to cars starting from built charging station i does not exceed the number of cars initially located at station i , provided that the station is built. These constraints prevent the over-allocation of trips relative to the resources available at each station. Constraints 3.7 ensure that for each potential charging station i , the sum of trips assigned to any car h that was initially assigned at potential station i should be less than or equal to the total number of vehicles L_i that were initially placed at that station. These constraints enforce that assigning trips to cars must consider the initial strategic decisions regarding vehicle allocation at each station. Constraints 3.8 impose restrictions on all nodes except i_0 , $i \in I$. For each node i_t , the sum of outgoing arcs including waiting arcs and final collection arcs, cannot exceed the capacity of the station. These constraints ensure that throughout the planning period, the number of EVs simultaneously parked at station i does not surpass its number of charging slots. Constraints 3.9 ensure that cars can access only the built charging stations. Flow conservation constraints 3.10 guarantee that each car's route corresponds to a path in the TLG in each scenario s . Link constraints 3.11 ensure that when a trip is assigned to a car, the flow in the network must occur on one travel arc corresponding to that trip. It links the vehicle's route with the specific trip it is supposed to complete. Constraints 3.12 mandate that each car h must fully charge its battery after completing the service.



4

Methodology

4. Methodology

4.1. Overview of the LMSP Framework




Figure 4.1.: Methodological Framework

4. Methodology

The proposed LMSP (LSTM-MLP-MDN Stochastic Programming) Framework is an innovative solution to the Charging Station Location Problem (CSLP) under demand uncertainty. Figure 4.1 illustrates the LMSP Framework and the workflow of the methodological framework used in this research.

The LMSP framework operates through the following key mechanisms:

1. Data Input: Historical traffic data and contextual data are collected.
2. Deep Learning Model (LSTM-MLP-MDN): This component learns from historical traffic patterns to predict future traffic demand. It captures temporal dependencies, processes contextual information, and generates probability distributions of future traffic demand. Section 4.2 offers a detailed explanation of it.
3. Demand Scenario Generation: The output from the deep learning model is used to generate multiple realistic demand scenarios, each representing a possible future demand pattern with an associated probability.
4. Two-Stage Stochastic Programming Model: This stage makes strategic decisions on charging station locations and initial vehicle allocations, followed by optimizing operational decisions based on the generated scenarios.
5. Optimization Process: The stochastic programming model uses the generated scenarios to find a robust solution, balancing the trade-off between infrastructure costs and operational efficiency across all scenarios.
6. Output: The LMSP Framework ultimately outputs the optimized locations for charging stations and the corresponding initial vehicle allocations.

4.2. Deep Learning Model for Traffic Demand Forecasting

4.2.1. LSTM-MLP-MDN Architecture

Figure 4.2 shows the architecture of LSTM-MLP-MDN, which consists of three parts: LSTM, MLP, and MDN. The LSTM is responsible for capturing temporal dependencies from historical traffic data, while the MLP processes non-sequential contextual features. The MDN outputs a probability distribution over possible future demands, allowing the model to represent demand uncertainty effectively.

4.2.2. Multilayer Perceptron (MLP)

The Multilayer Perceptron (MLP) is a feedforward type of artificial neural network containing fully connected neurons where the activation function is nonlinear. This research uses MLP to process non-sequential contextual data, such as weather information, traffic accident information, and public holiday information. By learning complex patterns from these diverse features, MLPs can serve as a good complement to LSTM networks to better predict total traffic demand.

In our MLP model, the input dimension is 11, incorporating various contextual information. There are two fully connected hidden layers. The first fully connected layers map the input




Figure 4.2.: LSTM-MLP-MDN Architecture

features to the first hidden layer, which includes dropout and uses the rectified linear unit (ReLU) as an activation function. The second fully connected layer maps the features from the hidden layer to the output layer, also employing the ReLU activation function. In the output layer, the dimension matches the hidden layer size of the LSTM model. Note that we use Kaiming uniform to initialize weights in these two hidden layers, as this method is particularly well-suited for layers using the ReLU activation function [Lyu et al., 2021]. It helps to prevent issues such as vanishing and exploding gradients by appropriately scaling the weights at the beginning of training.

4.2.3. Long Short-Term Memory (LSTM)

Long Short-Term Memory (LSTM) is a type of recurrent neural network (RNN) proposed by Hochreiter and Schmidhuber [1997] to solve the long-existing problems of exploding and vanishing gradients in RNN. Further research has determined that LSTM is particularly well-suited for time series prediction due to its proficiency in capturing long-short dependencies and temporal patterns in sequential data [Song et al., 2019]. This research leverages historical

4. Methodology

traffic demand data to predict future demand patterns, with a specific focus on predicting hourly traffic demand based on past observations. An LSTM unit (see figure 4.3) typically consists of a cell (the memory part of the LSTM) and three regulators of the cell's information flow, known as gates (input, output, and forget gate). A comprehensive explanation can be found in the appendix C.1.




Figure 4.3.: The mechanism of LSTM

The LSTM model in this research is a stacked 2-layer model designed to learn more complex patterns from data. The input data is a discrete sequence of hourly traffic demand from historical data (x_1, x_2, \dots, x_t) , where sequence length t is 24, representing a full day of hourly observations. This sequence is fed into the first layer at each time step t ($t = 1, 2, \dots, t$). Initializing the hidden state and cell state with zeros ensures that the LSTM starts with a neutral state for each input sequence. The first LSTM layer includes dropout to prevent overfitting, it acts as a feature transformer, converting 24 traffic demand sequences into higher-level representations and passing the hidden state h_t^1 to the second LSTM layer. The second LSTM layer can learn from this transformed data and outputs the hidden state h_t^2 of the last time step to the MDN model.

4.2.4. Mixture Density Network (MDN)

In this research, Mixture Density Network (MDN) is employed to handle the uncertainty and variability in traffic demand predictions. Unlike traditional models that output a single point estimate, the MDN outputs parameters for a mixture of Gaussian distributions, enabling the prediction of a full probability distribution over possible future traffic demands.

Our MDN model comprises three fully connected layers, each dedicated to parameterizing the components of a Gaussian mixture model, mixing with weight (or probabilities) π , standard deviations σ , and mean value μ . In the concatenation layer, the outputs of the MLP and LSTM are combined as input.

The first fully connected layer outputs the mixing weights π_i , utilizing a uniform initialization to promote balanced initial probabilities, thereby ensuring that no single component dominates the mixture prematurely. This layer employs a softmax activation function to normalize the weights, guaranteeing that they sum to 1 and thus conform to the probabilistic requirements of a mixture model. The second and third layers outputs the mean value μ_i and the standard deviations σ_i of Gaussian components, respectively. Both layers are initialized using the Kaiming normalization method and employ a softmax activation function to ensure positivity.



5

Case Study

5. Case Study

To validate and demonstrate the effectiveness of our proposed LMSP Framework, we will conduct a case study in Manhattan using the New York City (NYC) taxi dataset as a real-world dataset in this Chapter. Firstly, in section 5.1, we will perform data preprocessing on the NYC taxi dataset by selecting potential charging station locations and constructing the contextual data to support the prediction of total traffic demand. Subsequently, in section 5.2, we will apply the LSTM-MLP-MDN model to predict hourly total traffic demand and evaluate its performance against other potential models, using a set of evaluation metrics to confirm its superior predictive accuracy. In section 5.3, we will select five critical time periods on January 1, 2020, each representing a distinct traffic pattern. Using the LMSP Framework, the optimal deployment plan for charging stations will be derived based on these time periods. Following this, we will conduct a comparative evaluation to assess the added value of the LMSP Framework. Lastly, in section 5.4, we will provide a comprehensive analysis of the optimal charging station location deployment plan and offer recommendations for further improvements.

5.1. Dataset and Preprocessing



Figure 5.1.: Taxis of New York City

5.1.1. New York City (NYC) taxi dataset

New York City (NYC) is one of the largest and most densely populated cities in the world, and the NYC taxi open dataset includes a substantial amount of data rich in information over several years, making it a go-to resource for researchers [Shaygan et al., 2022]. The dataset used in this research is the NYC taxi data from 2015 to 2019, focusing on yellow and green taxi pickups and drop-offs within 25 selected popular regions in Manhattan. The reasons for choosing this specific dataset are as follows:

5. Case Study

1. **Comprehensive Time Span:** The five-year time span provides a comprehensive view of total traffic demand variations.
2. **Pre-Pandemic Data:** Data prior to 2020 is selected to avoid the significant impact of the COVID-19 pandemic, which drastically affected total traffic demand [TLC Factbook, 2024].
3. **Dominant Mode of Transportation:** During this period, taxis were still the primary mode of transportation in Manhattan. For-hire services like Uber and Lyft had not reached their current levels by 2023, making this dataset more representative of typical travel patterns from 2015 to 2019 [TLC Factbook, 2024].
4. **Representative Travel Behavior:** Taxi trips and car-sharing services exhibit similar travel behaviors, making the taxi data an appropriate substitute for studying car-sharing demand.




Figure 5.2.: Heatmap of Manhattan Traffic Demand from the Perspectives of Departure and Arrival
 Figure 5.2a and figure 5.2b respectively illustrate the annual average number of trips originating from and arriving at various regions from 2015 to 2019. Darker colors indicate a higher number of trips originating from or arriving at the regions.

Each trip in the NYC dataset contains the following information: origin point, destination point, starting time, and ending time. Both the origin and destination points are linked to a regional location ID. By counting the number of trips originating in each region, we can determine the traffic demand in different regions of Manhattan, as shown in Figure 5.2. Based on this analysis, we select 25 regions, focusing on areas with higher traffic (indicated by darker colors on the heat map) and key locations in Manhattan, such as Wall Street (see Figure 5.3). A detailed examination of these regions using Google Maps helps pinpoint specific coordinates suitable for establishing charging stations. The primary characteristics of each region are also identified, as summarized in Table A.1. The selected locations are

5. Case Study

strategically placed in popular areas, such as near hospitals, shopping malls, and notable landmarks.

Subsequently, we filter the Origin-Destination (OD) demand trips from the NYC taxi dataset, focusing on trips that begin and end within these 25 regions to represent the demand for the car-sharing service. Furthermore, we aggregate and individually analyze the hourly total traffic demand from 2015 to 2019 in these 25 regions. This data will serve as the basis for predicting hourly traffic demand in Section 5.2.

In addition to the above data processing steps, we need to address a critical component for our model's constraints: the travel time for each trip. This information is essential for implementing the battery limitation constraint 3.4 in the optimization model. To obtain accurate travel times, we employ Dijkstra's algorithm implemented through the NetworkX library [Hagberg et al., 2008]. This algorithm computes the minimal travel time between given origin and destination points. The speed data used in these calculations is sourced from Uber Movement Data 2019, providing realistic estimates of travel speeds in different parts of Manhattan at various times.

5. Case Study




Figure 5.3.: 25 Potential Locations of Charging Stations

5. Case Study

Table 5.1.: 25 Potential Locations of Charging Stations

Location ID	Latitude	Longitude	Region	Main Type
12	40.7045416	-74.0142564	Battery Park	Tourism
41	40.8075917	-73.9549547	Central Harlem	Residential
43	40.7815194	-73.9627249	Central Park	Tourism
48	40.7647531	-73.9880232	Clinton East	Mixed Use
50	40.7712271	-73.9935659	Clinton West	Commercial
68	40.7477804	-74.000438	East Chelsea	Commercial
74	40.80182	-73.9392953	East Harlem North	Residential
79	40.7260512	-73.9835308	East Village	Entertainment (Nightlife)
88	40.7034746	-74.0115948	Financial District South	Commercial
113	40.7346805	-73.9946012	Greenwich Village North	Educational
140	40.7611816	-73.9579018	Lenox Hill East	Residential
142	40.7738424	-73.9821234	Lincoln Square East	Residential
161	40.7585437	-73.9772064	Midtown Center	Commercial
163	40.7663007	-73.9818927	Midtown North	Commercial
170	40.7482478	-73.9762946	Murray Hill	Mixed Use
186	40.749101	-73.992006	Penn Station/Madison Sq West	Commercial
194	40.7841865	-73.9266152	Randalls Island	Entertainment (Sports)
231	40.7197847	-74.0068153	TriBeCa/Civic Center	Commercial
232	40.7152425	-73.9842337	Two Bridges/Seward Park	Residential
234	40.7379242	-73.9922478	UN/Turtle Bay South	Governmental
236	40.7801748	-73.9550942	Upper East Side North	Residential
237	40.76413	-73.9688047	Upper East Side South	Residential
239	40.7818426	-73.979274	Upper West Side South	Residential
243	40.856503	-73.932761	Washington Heights North	Residential
244	40.8369673	-73.9401365	Washington Heights South	Residential

5.1.2. Contextual Data of Manhattan

Some contextual factors are also considered in this model to help make more accurate traffic demand predictions. These data types include:

- **Weather Data:** Weather data includes maximum temperature, minimum temperature, mean temperature, precipitation (rain + snow), snow depth, and wind speed [Open Meteo, 2024].
- **Holiday Data and Big Events:** Public holidays such as Christmas Day, Easter Sunday, and big events like Presidents' Day can significantly affect traffic demand.
- **Traffic Accidents Data:** This dataset includes all motor vehicle collision counts recorded in the city of Manhattan [Department, 2024].

All contextual data need to be normalized before inputting them into the MLP model. This step is crucial because it helps train the neural network to converge faster and more efficiently. Here, holiday and big event data are digitized using One-Hot Encoding. The motor vehicle collision counts data are scaled with the StandardScaler because they follow a normal distribution, whereas all other contextual data are scaled using the MinMaxScaler.

5.2. Traffic Demand Prediction

5.2.1. Experiment Setup

The deep learning model LSTM-MLP-MDN is used to predict the probability distribution of hourly traffic demand in Manhattan. All training of the model was performed on a system with the following specifications: CPU model: AMD Ryzen 9 7945HX with Radeon Graphics, 2.50 GHz, and 16GB RAM. The specific parameters of the model are detailed in Table 5.2 below:

Table 5.2.: Hourly Traffic Demand Hyperparameters

Parameter	Value
Data Parameter	
seq_length	24
Model Parameters	
batch_size	32
epochs	15
learning_rate	0.001
num_gaussians	5
LSTM Parameters	
lstm_hidden_layer_size	88
dropout	0.4
MLP Parameters	
mlp_hidden_dim	5
Early Stopping	
patience	10
delta	0.05
Regularization	
L2 λ	0.1
entropy_weight	0.03

Sliding window cross-validation is utilized to train the deep learning model, as it helps in capturing temporal traffic demand patterns and prevents overfitting to particular time periods, ensuring the model generalizes effectively over five years. As shown in Figure 5.4, five windows are created to process all traffic demand data, with each window containing 90% training data and 10% testing data. A 30% overlap between adjacent windows is employed to capture more temporal dependencies and improve generalization to new and unseen data. By calculating the average loss across all windows, we can evaluate the model's performance. Note that the negative log-likelihood (NLL) is utilized as a loss function because it directly measures how well the predicted probability distributions fit the observed data, accommodating the inherent uncertainty in traffic demand predictions.

Furthermore, an early stopping mechanism is introduced into the model. This is necessary because the model is trained on datasets from five different windows. Without proper control, training within a single window can lead to overfitting, which subsequently diminishes performance in subsequent windows. The early stopping mechanism can help in obtaining

5. Case Study



Figure 5.4.: Sliding window cross-validation

the best model by monitoring performance and stopping at the point of optimal performance rather than continuing to train and potentially degrading the model's performance. In Table 5.2, "patience" refers to the number of consecutive epochs allowed without an improvement in the validation loss before stopping the training process, and "delta" is the minimum change in the validation loss that is considered a significant improvement.

Regularization is also crucial in this to ensure the model utilizes all the mixture components effectively rather than concentrating on just one or two. Here, we utilize two regularization strategies: L2 regularization (weight decay) and entropy regularization. L2 regularization can penalize large weights in the MDN network, which helps in preventing overfitting. In Table 5.2, 'L2 λ ' is the regularization parameter for L2 regularization, which is added to the loss function as $\lambda \sum w_i^2$, where w_i are the weights of the MDN network. On the other hand, entropy regularization can encourage the MDN network to use all components in the Gaussian mixture more uniformly. In table 5.2, entropy_weight is utilized to control the importance of entropy regularization. This term is added to the loss function as $\alpha \sum \pi_i \log(\pi_i)$, where π_i are the mixture component weights and α is the entropy weight.




Figure 5.5.: Training and Validation Loss over Five Windows

Figure 5.5 shows the training results for the model over 5 different windows. Both training

5. Case Study

and validation losses consistently and steadily decrease across all epochs and windows, indicating that the model learns and fits the training data well. In each window, the validation loss decreases almost in parallel with the training loss, indicating that the model is unlikely to overfit. Furthermore, the final loss values for both training and validation are quite close to each other, and both are significantly lower than the initial values. This suggests that the model's performance on the validation data is nearly as good as on the training data, indicating good generalization.

5.2.2. Evaluation Metric

According to the literature review in Section 2, various evaluation metrics are used to assess model performance. Since the model's output is a probability distribution over traffic demand rather than a single value, the evaluation metrics need to account for this. This research evaluates the model from the following aspects:

Model Prediction Accuracy

To evaluate the model's prediction accuracy, the most likely predicted value from the probability distribution is considered because it represents the single most probable outcome, making it straightforward and intuitive. In the LSTM-MLP-MDN model, it is the mean value of the normal distribution which holds the highest weight. This predicted value is then compared to the actual observed value using metrics such as MAE (Mean Absolute Error) and MAPE (Mean Absolute Percentage Error), which are defined as:

$$\text{MAE} = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i| \quad (5.1)$$

$$\text{MAPE} = \frac{1}{n} \sum_{i=1}^n \left| \frac{y_i - \hat{y}_i}{y_i} \right| \times 100 \quad (5.2)$$

where n is the number of observations, y_i is the actual traffic demand value for the i -th observation, and \hat{y}_i is the predicted traffic demand for the i -th observation. MAE provides a clear and interpretable measure of average prediction error that is robust to outliers, while MAPE allows for comparison of forecast accuracy across datasets with different scales, and its percentage format is intuitive and easy to understand.

R-squared (R^2)

The R^2 metric measures the proportion of the variance in the dependent variable that is predictable from the independent variables. It provides an indication of how well the predicted values approximate the actual data points. A higher R^2 value indicates a better fit of the model to the observed data.

$$R^2 = 1 - \frac{SS_{res}}{SS_{tot}} \quad (5.3)$$

5. Case Study

where:

SS_{res} (Sum of Squares of Residuals) is the sum of the squared differences between the observed values and the predicted values by the model. SS_{tot} (Total Sum of Squares) is the sum of the squared differences between the observed values and the mean of the observed values.

Prediction Interval Coverage Probability (PICP)

The Prediction Interval Coverage Probability (PICP) is used to evaluate the reliability of the model's probabilistic predictions. It measures the proportion of times the actual traffic demand values fall within the predicted 95% confidence interval. A high PICP indicates that the model's predicted intervals effectively capture the true variability in the data.

The 95% confidence interval in the LSTM-MLP-MDN model can be derived by sampling as follows:

- **Sample Generation:** Draw 100 samples from the Gaussian mixture distribution. This sample size is sufficient to ensure statistical stability.
- **Sorting:** Sort the samples to create an empirical distribution.
- **Percentile Calculation:** Determine the 2.5th percentile and the 97.5th percentile of the sorted samples.
- **Prediction Interval:** The 95% prediction interval is defined by the 2.5th percentile and the 97.5th percentile of the sorted samples.

Log-Likelihood (LL)

Log-likelihood (LL) is used to evaluate the goodness-of-fit of the model on the validation set. It measures how well the predicted probability distributions match the observed traffic demand data. A higher LL value indicates a better fit, as it implies that the model assigns higher probabilities to the observed data points.

In the LSTM-MLP-MDN model, since the output is a mixture of normal distributions, equations C.7 and C.8 illustrate the process to calculate the mixture probability distribution. The Log-Likelihood is then calculated as follows: for each validation sample in the validation dataset, the log of the mixture probability distribution function is computed, which is denoted as:

$$\log \mathcal{L}(y) = \log \left(\sum_{i=1}^N \pi_i f_i(y) \right) \quad (5.4)$$

For the entire validation dataset, sum these log-likelihood values:

$$\log \mathcal{L}_{\text{total}} = \sum_{y \in \text{validation set}} \log \left(\sum_{i=1}^N \pi_i f_i(y) \right) \quad (5.5)$$

5.2.3. Comparison With State-of-the-Art

This section compares our proposed LSTM-MLP-MDN model with several well-established models in the field of time series prediction and traffic demand forecasting. These models are chosen for their proven performance in similar contexts and their ability to handle complex time-dependent data. The comparison aims to demonstrate the advantage of our approach in terms of accuracy and ability to capture uncertainty in traffic demand predictions.

Our proposed model is compared with the following potential models. Note that all parameters of these models have been carefully adjusted through rigorous tuning processes to ensure optimal performance, allowing for a fair comparison.

LSTM-MDN

Compared to the previous models, the LSTM-MDN model has removed the MDN's component responsible for extracting contextual information. After multiple rounds of hyperparameter tuning, the optimal parameters are as follows:

seq_length = 24, *batch_size* = 32, *epochs* = 15, *learning_rate* = 0.001, *num_gaussians* = 5, *lstm_hidden_layer_size* = 88, *dropout* = 0.3.

Bootstrap Method

While LSTM-MDN naturally produces probability distributions, traditional time series models like Auto-Regressive Integrated Moving Average (ARIMA) and Seasonal Auto-Regressive Integrated Moving Average (SARIMA) are designed to produce point estimates. To ensure an appropriate comparison with our proposed LSTM-MLP-MDN model, we employ the Bootstrap method for both ARIMA and SARIMA models to generate probability distributions.

The Bootstrap method is a resampling technique that allows us to estimate the sampling distribution of an estimator by resampling with replacement from the original sample. In the context of our time series models, it helps convert point estimates into probability distributions, thus providing a measure of prediction uncertainty. The process involves resampling the residuals from the original time series predictions to create multiple simulated samples. Each resampled series is used to generate new forecasts, and this process is repeated 1000 times to create an empirical distribution of predictions. The detailed method integrating ARIMA, SARIMA, and Bootstrap is as follows:

- **ARIMA - Bootstrap:** The ARIMA model combines autoregressive and moving average components and includes differencing to make the time series stationary. For hourly prediction, the ARIMA model parameters are set as $p = 2$, $d = 0$, and $q = 3$. To obtain a probability distribution as output, the Bootstrap method is applied.
- **SARIMA - Bootstrap:** The SARIMA model extends ARIMA by incorporating seasonal components. It is well-suited for traffic demand with seasonal patterns, showing daily and weekly trends. For hourly prediction, parameters are set as order (3,0,5) and seasonal order (1,1,1,24). For daily prediction, parameters are set as order (1,1,1)

5. Case Study

and seasonal order (2,1,2,7) where the Akaike Information Criterion is minimized. Similarly, the Bootstrap method is also used.

5.2.4. Results Analysis

Model Performance Comparison

Table 5.3.: Comparison of the model results for hourly prediction

Experimental model	MAE	MAPE (%)	R^2	PICP (%)	LL
LSTM-MLP-MDN	174.3023	15.10	0.8872	74.60	-23.0259
LSTM-MDN	179.2215	14.73	0.8592	65.82	-20.8965
ARIMA-Bootstrap	306.4000	38.3892	0.09	9.01	-17425.8749
SARIMA-Bootstrap	263.2236	26.69	0.12	12.3	-16823.69

Table 5.3 presents a comprehensive comparison of the performance metrics for the models mentioned in the previous section in hourly traffic demand prediction. It is evident that the LSTM-MLP-MDN model outperforms the others. Specifically, the Mean Absolute Error (MAE) is 174.3023. Given that the hourly traffic demand in historical data ranges around 2500, this value is acceptable. The Mean Absolute Percentage Error (MAPE) value of 15.10% indicates the average percentage error between the predicted and actual values. A MAPE of 15.10% is relatively reasonable for traffic prediction, considering that traffic patterns can be highly volatile and influenced by numerous external factors. An R^2 value of 0.8872 indicates that approximately 88.72% of the variance in the traffic demand is explained by the model. This high value, significantly superior to other models, demonstrates the model's strong explanatory power and excellent fit to the data. The Prediction Interval Coverage Probability (PICP) is 74.60% which is the highest value compared to other models. This indicates that the model's predicted intervals capture the true values more effectively, showcasing its superior ability to quantify uncertainty in predictions. The log likelihood (LL) value is -23.0259, which is higher than that of other traditional methods.

To conclude, the superior performance of the LSTM-MLP-MDN model across most metrics, particularly in MAE, R^2 , and PICP, underscores its effectiveness in predicting traffic demand probability distributions. Its ability to capture both the central tendency (as evidenced by low MAE and high R^2) and the uncertainty (high PICP) of traffic demand makes it the most suitable choice for generating reliable inputs for subsequent stochastic programming optimization models.

Hourly Traffic Demand Pattern and Prediction Validation

Figure 5.6 provides a comparison of total traffic demand on January 1st from 2015 to 2020 in 25 selected regions in Manhattan. This visualization provides valuable insights into historical traffic patterns:

- **Hourly Traffic Demand Pattern:** The graph clearly delineates typical daily traffic patterns, including peak and non-peak hours. Consistently across the years, we observe

5. Case Study




Figure 5.6.: Total Traffic Demand on January 1st by Year (2015-2020), Predicted for January 1st, 2020

that traffic starts at a low level early in the morning, rises sharply to peak around midday, and then decreases towards the evening. This recurring pattern not only demonstrates the stability of Manhattan’s daily traffic rhythm but also provides crucial insights for our subsequent analysis and decision-making process. By understanding these fluctuations, we can strategically select representative time periods (see Section 5.3.1), accounting for various future traffic patterns in Manhattan, ranging from periods of low demand to high demand.

- **Long-term Trend Analysis:** A significant trend observed is the steady decrease in total traffic demand from 2015 to 2020. According to [TLC Factbook \[2024\]](#), one major factor contributing to this trend is likely the rapid rise of for-hire vehicle services such as Uber and Lyft, which have steadily captured market share from traditional taxis since 2015. In the long term, future traffic demand is expected to decrease. However, since the data used here is based on NYC taxi trips, there are still differences compared to the potential demand for electric car-sharing services in the future. Therefore, we cannot directly equate this downward trend with changes in demand for electric vehicle-sharing services. Instead, our analysis focuses on short-term demand patterns, specifically hourly patterns. These short-term patterns are likely to have a higher similarity between taxi usage and the potential demand for electric vehicle-sharing services.

The LSTM-MLP-MDN model demonstrates its ability to capture the main traffic trends throughout the day. While there are some deviations in absolute values, these are within an acceptable range considering the complexity of hourly predictions and the potential for noise and fluctuations in such granular data. This lays a strong foundation for subsequent analysis and planning.

5. Case Study




Figure 5.7.: Predicted Total traffic demand in Manhattan on January 1, 2020

5.3. Charging Station Deployment Plan

After completing the hourly traffic demand prediction, we focus on January 1, 2020, to demonstrate the practical application of our model. This date is chosen as it represents the most recent data point in our dataset, allowing us to showcase the model’s predictive capabilities on current traffic patterns.

Our LSTM-MLP-MDN model generates a probability distribution of traffic demand for each hour, specifically providing five potential demand levels with corresponding probabilities. We then integrate these probability distributions as scenarios into our stochastic programming model. This integration allows us to obtain an optimal EV charging station deployment plan considering demand uncertainty within each hour.

To provide a comprehensive analysis and actionable recommendations for the deployment plan in section 5.4, we consider Manhattan’s diverse traffic demand patterns throughout the day. We strategically select five specific time periods that represent distinct traffic conditions (see section 5.3.1). For each time period, we obtain a distinct deployment plan from our proposed LMSP Framework, allowing us to understand how the optimal infrastructure setup changes based on time-specific demand patterns.

Furthermore, to demonstrate the superiority of the LMSP Framework, we conduct a comparative analysis (see section 5.3.5). We evaluate the added value of our approach against traditional methods combined with stochastic programming. This comparison specifically focuses on how the deep learning component of our framework provides more accurate probability distributions of traffic demand. It aims to illustrate that the improved quality of input data from the deep learning model enables the stochastic programming component to make more informed decisions, ultimately resulting in more efficient and cost-effective charging station deployment strategies.

5.3.1. Determining Critical Time Periods

Figure 5.7 shows predicting total traffic demand in Manhattan City on January 1, 2020. Actually there is a significant fluctuation in traffic patterns throughout each day. Therefore,

5. Case Study

to further understand the optimal charging station deployment under the different traffic patterns, we divide the day into seven segments based on possible traffic conditions and select five specific time periods (points marked in red in figure 5.7) as five representative time periods for detailed analysis:

- **Time period 1:** All the requested trips start at 4h. This low-traffic period may include evening entertainment activities, some night shift workers, and urgent travel needs.
- **Time period 2:** All the requested trips start at 9h. This period has high commuter traffic due to people going to work and school. Major routes include travel from residential neighborhoods such as the Upper West Side and Harlem to business districts like Midtown Manhattan and the Financial District [TLC Factbook, 2024].
- **Time period 3:** All the requested trips start at 15h. This non-peak period with shopping, entertainment, and business activities.
- **Time period 4:** All the requested trips start at 18h. This time period also has high commuter traffic due to people returning home from work and school, along with evening dining out. Key routes include trips from Midtown Manhattan to outer boroughs and residential neighborhoods. Notable patterns include travel from the Financial District to Brooklyn Heights and from Times Square to Queens [City of New York, 2024].
- **Time period 5:** All the requested trips start at 21h. This time period has moderate traffic due to entertainment and social activities, with some late-night shopping. Popular routes include travel between entertainment hubs like Times Square and residential areas in the Upper West Side, as well as from Greenwich Village to the East Village TLC Factbook [2024].

5.3.2. Assigning Predicted Traffic Demand to Specific Routes

After predicting the total traffic demand, we use the frequency statistic from historical data (2015 to 2019) to determine the probability of specific traffic demand for each route, where both the origin point and destination point are among the selected 25 points. This allows us to assign total traffic demand to each specific route based on their respective probabilities. The specific data can be obtained by sampling from similar historical data.

For example, when collecting the traffic demand data starting at 9 AM on January 1, 2020, before inputting it into the stochastic programming model, we first obtain the predicted total traffic demand for 9 AM, which is 2687. Then, we assign this total traffic demand to each route based on their probability distribution using historical data. Specifically, we filter out all trips with a starting time of 9 AM from January 1 of each year from 2015 to 2019. After removing the year information and shuffling the data, we sample the trips the required number of times to represent the requested trips for each specific route at 9 AM on January 1, 2020.

5.3.3. Traffic Demand Level Scenario Sampling

LSTM-MLP-MDN outputs a mixture of Gaussian distributions as traffic demand predictions for each particular hour. Since the two-stage stochastic programming model requires different demand levels with corresponding probabilities, we sample five different values

5. Case Study

Table 5.4.: Probabilistic Distribution of Predicted Traffic Demand Over Five Time Periods

(a) 4h Prediction Result		(b) 9h Prediction Result		(c) 15h Prediction Result	
Demand	Probability	Demand	Probability	Demand	Probability
116	0.874495	2687	0.517256	2968	0.479632
172	0.122543	2582	0.436146	3016	0.477480
175	0.002473	1919	0.027194	2194	0.023855
537	0.000367	2756	0.014608	3075	0.014317
1189	0.000122	2708	0.004795	2956	0.004716

(d) 18h Prediction Result		(e) 21h Prediction Result	
Demand	Probability	Demand	Probability
3433	0.528412	3285	0.539734
3120	0.413350	3042	0.399869
2449	0.030587	2379	0.031845
3213	0.020709	3116	0.021401
3063	0.006943	3002	0.007152

from the mixture distribution to represent these varying demand levels. The specific sample process below can ensure that select rapidly 5 different level values with corresponding probabilities:

1. Extract traffic demand values (mean value), probabilities (weight), and standard deviation from the prediction results .
2. Calculate the range of integer values to evaluate demand. Here the range is denoted as:

$$[D_{min} - 3\sigma_{max}, D_{max} - 3\sigma_{max}] \quad (5.6)$$

where:

D_{min} and D_{max} represent the minimum mean value and the maximum mean value of the mixed Gaussian distribution respectively. And σ_{max} represents the minimum standard deviation.

3. For each demand value, calculate the Gaussian distribution probability density function (PDF) for each integer in the range and accumulate these probabilities weighted by probabilities in step 1.
4. Normalize the cumulative probabilities to ensure their sum is 1, converting the cumulative probability values into a valid probability distribution.
5. Identify and store the top 5 most probable demand levels based on the cumulative probabilities. Results are shown in Table 5.4.

5. Case Study

Table 5.5.: Model Parameters

Parameter	Value	Meaning
$ S $	5	Number of Scenarios
$ I $	25	Number of potential charging station locations
T_{max}	60	Number of time intervals within an hour
f_i	3.42 (150000/ Δ)	Fixed cost of the station if built at location i
c	0.34 (15000/ Δ)	Purchasing cost of each car
i	20	Income from each trip
C	5	Capacity of station i
W	77 (3372600/ Δ)	Budget
i_k	20	Income from the k -th trip
N	30	Highest number of vehicles that can be purchased
Φ	55	Maximum travel time of a fully charged car
τ	15	Required charging time after a car is returned to the charging station

Table 5.6.: Gurobi Setup

Parameter	Value	Description
NodefileStart	0.5	Start writing node files to disk after 0.5GB of memory is used.
MIPFocus	2	Focus on finding feasible solutions quickly.
Cuts	3	Aggressive cut generation, speeding up solving and improving solution quality.
MIPGap	0.05	The solver will stop only when the gap between the current solution and the theoretically optimal solution is less than 5%.
Presolve	2	Aggressive resolve to simplify the model and reduce solving time.
Heuristics	0.5	Helps in quickly finding good feasible solutions.
Threads	32	Utilizes all logical processors.

5.3.4. Optimal Charging Station Deployment Results

This section respectively inputs the predicted hourly traffic demand probability distributions of these five time periods into the optimization model of section 3.2 to obtain the final optimal deployment plan. According to the 2022 Annual Report of the NYC Taxi & Limousine Commission, considering millions of trips in Manhattan, the high-speed charging requirement, and the high-frequency usage, the proposed charging stations are DC fast charger stations. Parameters such as the cost of a DC fast charger station (approximately 150000 dollars) and the purchasing cost of an EV (15000 dollars) are also provided in the report [New York City Taxi & Limousine Commission, 2022]. Assuming that the service life of all charging stations and purchased EVs is 5 years, cost parameters should be adjusted by dividing the time factor Δ (where $\Delta = 365 \times 24 \times 5$) to convert them into hourly coefficients. These adjustments allow the use of hourly parameters in the optimization model, thereby accurately reflecting the cost within the hourly demand scenario. The model parameters are set in Table 5.5. Considering the complexity of the optimization model, which includes a substantial number of constraints and variables, it is necessary to establish specific parameters to improve both the speed and quality of the solution. The specific parameters are detailed in Table 5.6.

Our optimization process produced distinct deployment plans for each of the five time periods studied: 4h, 9h, 15h, 18h, and 21h. These plans reveal variations in the number and location of charging stations, as well as the allocation of vehicles, reflecting the changing demand patterns throughout the day. The detailed deployment plans for each time period are provided in appendix A for reference. A comprehensive analysis of these deployment plans, including temporal and spatial considerations, will be presented in section 5.4.

5.3.5. The Added Value of Deep Learning in Stochastic Programming for Demand Prediction

Traditional Approaches to Comparison

To highlight the benefits of incorporating a deep learning component of our proposed LMSP Framework, we compare its performance against two traditional methods combined with stochastic programming. These two traditional methods rely on frequency statistics from historical data, as detailed below.

Historical Frequency-based Approach (HFA): This method utilizes historical data from 2015 to 2019, categorizing hourly total traffic demand into 24 classes based on the hour of the day. For each hourly class, we select the top 5 most frequent total traffic demand values. The probability of each value is calculated as its frequency within the class. These probabilities are then normalized to ensure they sum to one. This approach captures the most common demand patterns but may not adequately represent the full range of demand variability.

Gaussian Distribution Approach (GDA): Similar to the Historical Frequency-based Approach, this approach first categorizes historical data into 24 hourly classes. For each class, a Gaussian distribution is fitted to the traffic demand data. Five representative values are extracted from this distribution at the 0.2, 0.35, 0.5, 0.65, and 0.8 quantiles. These values are rounded to integers to represent discrete demand levels. Finally, the probabilities are

5. Case Study

normalized to sum to one. This method provides a more nuanced representation of demand variability compared to the Historical Frequency-based Approach, as it considers the underlying distribution of the data.

Key Performance Indicators (KPIs)

To quantitatively evaluate the strengths and weaknesses of these approaches, we employ the following key performance indicators (KPIs). Note that the specific meaning of the symbols has been defined in detail in Section 3.2.

Profit: This is the objective value of the optimization model in section 3.2.3. It reflects the total profit obtained by the system through the optimal allocation of charging stations and vehicles.

Return on Investment (ROI): This measures the efficiency of the system's capital investment, defined as the ratio of the profit gained to the initial investment, which specifically includes the sum of the charging station construction costs and the EV purchasing cost.

$$\text{ROI} = \frac{\text{Profit}}{\sum_{i \in I} f_i y_i + c \sum_{i \in I} L_i} \quad (5.7)$$

Demand Satisfaction Ratio (DSR): This KPI can provide a comprehensive measure of how well the car-sharing system meets demand across all scenarios. It is calculated as follows:

$$\text{DSR} = \sum_{s \in S} p_s \frac{\sum_{k \in K^s} x_k}{|K^s|} \quad (5.8)$$

where,

- $|K^s|$ is the total number of requested trips in scenario s .

Charging Station Utilization Rate (CSU): Firstly, the utilization rate for charging station i at time t is defined as follows:

$$\text{UtilizationRate}_i^t = \frac{\sum_{h \in H} f_a^h}{C}, \quad \forall a \in A_W^s(i, t), \quad \forall s \in S \quad (5.9)$$

where,

- $A_W^s(i, t)$ represents the set of waiting arcs at station i during time t in scenario s .
- C is the capacity of each charging station, i.e., the number of charging slots available at each charging station.

Then, the average Utilization rate over the planning period is calculated by taking the mean of the utilization rates at the T time points within the planning period. In this research, the planning period is 1 hour, with each time point representing one of the 60 minutes within that hour.

$$\text{AvgUtilization}_i = \frac{1}{T} \sum_{t=1}^T \frac{\sum_{h \in H} f_a^h}{C}, \quad \forall a \in A_W^s(i, t), \quad \forall s \in S \quad (5.10)$$

5. Case Study




Figure 5.8.: Performance Comparison of Different Methods Over Time Using KPIs

Finally, weighted utilization for different scenarios (denoted as CSU) is:

$$CSU = \sum_{s \in S} p_s \left(\frac{1}{T} \sum_{t=1}^T \frac{\sum_{h \in H} f_a^h}{C} \right), \quad \forall a \in A_W^s(i, t), \quad \forall s \in S \quad (5.11)$$

This KPI can measure how well the charging infrastructure is being used.

Results Analysis: Traditional-based Stochastic Programming vs. LMSP Framework

Through analysis of Figure 5.8, we can evaluate the performance of the LMSP Framework from the following aspects. The specific data is provided in Appendix D.

Profit and ROI: In all time periods, traditional methods HFA and GDA show higher profit and ROI compared to the LMSP Framework. Further analysis reveals that this is mainly due to the higher demand prediction levels of HFA and GDA. Given the short planning period (1 hour) in this research, and the optimization model parameters setting where income from all car-sharing services far exceeds costs, demand levels significantly impact profit and ROI. However, the high demand predictions of HFA and GDA may lead to overly aggressive resource allocation. While this results in higher short-term profits, it may be accompanied by unnecessary resource waste, such as the over-allocation of charging stations and vehicles.

Demand Satisfaction Ratio (DSR): The LMSP framework consistently demonstrates significantly higher DSR compared to the HFA and GDA methods in all time periods. This indicates that the LMSP framework can better meet actual demand at various time points. While HFA and GDA's high demand predictions lead to greater resource deployment, they

do not accurately match real demand, resulting in lower satisfaction rates for user requests. The LMSP framework, through more precise prediction of demand probability distributions, enables the model to allocate resources more closely aligned with actual demand, avoiding over- or under-allocation. Consequently, the LMSP framework can better satisfy user needs, which is particularly crucial in electric car-sharing systems where user satisfaction is directly related to service quality.

Charging Station Utilization (CSU): The LMSP Framework shows notably superior performance in CSU compared to other methods, indicating its ability to better utilize charging station resources across different time periods. HFA and GDA, due to their higher demand prediction levels, may lead to the over-allocation of charging station resources, which are not actually fully utilized. The LMSP framework, through more accurate demand prediction, allocates charging station resources more rationally. This ensures effective use across all time periods, avoiding resource waste or idleness.

Conclusion:

Although HFA and GDA perform better in terms of profit and return on investment, this is mainly due to short-term gains from their high-demand predictions. However, the LMSP Framework, through precise demand prediction using deep learning, achieves significantly better demand satisfaction rates and charging station utilization rates. Specifically, compared to HFA, the LMSP Framework shows an average improvement of 26.32% in DSR and 36.17% in CSU across five time periods. When compared to GDA, it demonstrates an average increase of 14.55% in DSR and 12.06% in CSU. These substantial improvements in operational metrics indicate its ability to allocate resources more rationally, avoid waste, and perform superiorly regarding user satisfaction and resource utilization efficiency. Therefore, the LMSP Framework demonstrates clear advantages in addressing actual demand variations, improving system operational efficiency, and optimizing resource utilization. This suggests that the application of deep learning in stochastic programming indeed brings significant added value to the planning and operation of electric car-sharing systems.

5.4. Comprehensive Analysis for Charging Station Location Optimization

5.4.1. Overall Demand Patterns

To effectively evaluate the charging station deployment plan for the electric car-sharing system, it is essential to first develop a comprehensive understanding of Manhattan's traffic demand patterns. Our analysis focuses on two key dimensions: temporal demand patterns and spatial demand patterns.

Temporal Demand Patterns

The variation in traffic demand across different time periods is substantial, reflecting the diverse traffic patterns throughout the day. According to Figure 5.6, the total traffic demand across five specific time periods in 25 selected locations, distinct temporal patterns emerge:

5. Case Study

1. **Late Night (4h):** The lowest traffic demand is observed at this time, reflecting minimal travel activity during late night hours. This period likely serves to night shift workers, late-night entertainment, and emergency travel needs.
2. **Morning Peak (9h) and Evening Peak (18h):** These periods show the highest traffic demand, typical of rush hour patterns. The 9h peak corresponds to the morning commute, with people traveling to work and schools. The 18h peak reflects the evening rush, as commuters head home and participate in post-work activities.
3. **Afternoon (15h) and Night (21h):** These periods show moderately high demand, though not as intense as the peak hours. The 15h demand likely represents a mix of business travel, shopping, and leisure activities. The 21h demand might be attributed to evening social activities, late-night shopping, and entertainment.

Contrary to common belief, another interesting finding is that the morning and evening peak hours are not entirely symmetrical. It is often assumed that people commute from home to work during the morning peak and return home during the evening peak. However, our data shows an imbalance between these two periods. This difference could be due to factors like staggered work schedules, varying return trip patterns, or evening activities unrelated to work. This asymmetry highlights the need to adjust deployment strategies based on the unique timing patterns in each area.

5.4.2. Spatial Demand Patterns

The traffic demand in Manhattan exhibits significant spatial variability, reflecting the city's diverse and complex urban landscape. To gain insights into these patterns, we analyze historical data, aggregating trips originating from each region at consistent time points (specifically, January 1st) across multiple years. Figure 5.9 illustrates the resulting average number of trips, revealing distinct regional demand patterns. The data here is also used later in the figures, from Figure 5.10 to Figure 5.14 presented in the form of heat maps. For detailed data, refer to appendix B.




Figure 5.9.: Trips by starting region and hour

High-Demand Regions

Based on historical traffic data, several regions exhibit significantly higher traffic demand, especially during peak hours on weekdays, such as 9h and 18h. These regions are typically concentrated in the core residential and commercial areas of Manhattan. The following are the two areas with the highest demand.

Upper East Side South (Region 237): The Upper East Side South is a major high-income residential area, with traffic demand peaking during the morning and evening rush hours on weekdays. Historical data indicates that this region experiences significant commuting demand during these times. The high population density and many residents commuting to commercial areas in Manhattan contribute to these peak periods. Given the area's dense residential nature, prioritizing the construction of charging stations in this region is essential.

Midtown Center (Region 161): Midtown Center is Manhattan's principal commercial district, home to numerous office buildings and commercial establishments. As a result, traffic demand in this area peaks during weekday mornings (around 9 h) and evenings (around 18h). The demand pattern in this commercial region reflects the high volume of commuters working in the area. A well-distributed network of charging stations in Midtown Center is essential for supporting the heavy use of electric vehicles during these peak hours. This demand pattern suggests that priority should be given to expanding charging infrastructure in commercial districts to accommodate the needs of commuters and business-related travel.

Low-Demand Regions:

In contrast to high-demand regions, several areas in Manhattan show relatively low traffic demand. These regions are often associated with recreational or tourist activities, or they may be residential areas with specific characteristics. The following are examples of typical low-demand regions.

Randalls Island (Region 194): Randalls Island is primarily a recreational and sports area, with traffic demand concentrated around events and activities. Figure 5.9 shows minimal trips starting from this area across all time periods. Therefore, the allocation of charging station resources in this context requires careful consideration.

Two Bridges/Seward Park (Region 232): Although Two Bridges/Seward Park is a residential area, its traffic demand is significantly lower compared to high-demand residential regions like the Upper East Side. This may be attributed to a smaller population and fewer direct connections to major transit hubs or commercial centers. Given the lower commuting demand, extensive charging infrastructure may not be necessary in this area.

Relationship Between Demand and Region Type

When considered combined with table A.1, categorizes regions by their primary function (e.g., residential, commercial, mixed-use), we can gain deeper insights into the relationship between region type and travel behavior.

Residential Areas: Regions like Upper East Side South (Region 237), Upper East Side North (Region 236), and Upper West Side South (Region 239) display strong morning and evening

5. Case Study

peaks, consistent with typical commuting patterns. This is likely due to the high population density in these areas and the tendency for residents to commute to commercial districts for work. However, not all residential areas follow this trend, as seen in the lower demand in Washington Heights North (Region 243) and Washington Heights South (Region 244). This discrepancy may be attributed to factors such as lower population density, different socioeconomic characteristics, or better access to public transportation options.

Commercial Areas: Regions like Midtown Center (Region 161), Midtown North (Region 163), and the Financial District South (Region 88) show high demand during business hours, peaking in the afternoon and early evening. This pattern reflects the influx of workers and business-related travel, as well as potential lunchtime activities. The concentration of office buildings, businesses, and services in these areas drives the high daytime demand. These areas might benefit from a higher density of charging stations to accommodate the daytime influx of electric vehicles.

Mixed-Use Areas: In areas such as Clinton East (Region 48) and Murray Hill (Region 170), demand remains consistently high throughout the day. This steady demand is likely due to a balanced mix of residential and commercial activities, creating a more evenly distributed travel pattern throughout the day. These areas might require a more uniform distribution of charging stations to cater to both residential and commercial needs.

Entertainment Areas: The East Village (Region 79), known for its nightlife, sees increasing demand for the evening and night. This pattern reflects the area's concentration of restaurants, bars, and entertainment venues that attract visitors later in the day. In contrast, Randalls Island (Region 194), categorized as an entertainment area focused on sports, exhibits very low demand across all time periods. This could be due to its isolated location and event-driven nature, suggesting that charging infrastructure here might only be necessary during specific events.

Tourism Areas: Central Park (Region 43) has substantial demand, particularly in the afternoon, likely due to its popularity among both tourists and locals for recreational activities. Battery Park (Region 12), while also a tourist area, experiences comparatively lower demand. This difference might be attributed to Central Park's larger size, central location, and broader appeal for various activities throughout the day.

Educational Areas: Greenwich Village North (Region 113), which houses several educational institutions, shows demand peaking in the morning and remaining high throughout the day. This pattern likely reflects the daily routines of students and staff, including morning arrivals, daytime activities, and evening departures or events.

Governmental Areas: UN/Turtle Bay South (Region 234) experiences peak demand during standard business hours, particularly in the afternoon and early evening. This pattern reflects the concentration of governmental and diplomatic activity in the area, with demand driven by employees, visitors, and related business traffic.

5.4.3. Evaluation of Deployment Plan

We obtained optimal deployment plans based on five specific time periods from January 1, 2020. These plans identify the optimal locations for charging stations and the strategic allocation of vehicle resources. In this section, we will individually analyze the deployment strategies for each time period, concluding with an overall evaluation.

Time Period 1 (4h) (see Figure 5.10)

The deployment plan for the 4h time period reflects the unique characteristics of late-night traffic patterns in Manhattan. During this period, overall traffic demand is at its lowest, yet the plan demonstrates a strategic approach to resource allocation.

At 4h, charging stations are spread across a wide area, with 19 out of 25 possible locations having stations installed. Even though overall demand is low at this time, this wide coverage seems aimed at keeping the system accessible. Important high-demand areas, like Midtown Center (Region 161) and Upper East Side South (Region 237), have charging stations, which matches the steady demand seen in these areas throughout the day.

The distribution of vehicles is also notable. A total of 29 vehicles are placed strategically across the city, with the largest numbers assigned to certain key areas:

- Clinton East (Region 48): Clinton East (Region 48) receives the highest allocation of 4 vehicles, which fits with our earlier observation that mixed-use areas tend to have steady demand throughout the day. This large allocation likely anticipates both late-night returns and early-morning departures in this diverse urban area.
- Penn Station/Madison Sq West (Region 186): It has been assigned 5 vehicles, the highest allocation observed. This substantial commitment reflects the importance of major transportation hubs, even during off-peak hours, serving late-night arrivals and early-morning departures.
- Midtown Center (Region 161): It received 4 vehicles, recognizing its role as Manhattan's main commercial district. Despite the generally low commercial activity at this hour, this allocation reflects foresight in preparing for early-morning business operations and accommodating night-shift workers, as highlighted in our analysis of demand patterns.

These high-vehicle allocations in key areas contrast with the more moderate distributions elsewhere. For instance, East Village (Region 79) and Upper East Side South (Region 237) each receive 2 vehicles, reflecting their roles as residential and entertainment hubs with potential late-night or early-morning demand.

Interestingly, several stations, such as Central Harlem (Region 41), Central Park (Region 43), and East Harlem North (Region 74), have charging stations but no vehicles allocated. This strategy aligns with our observation that these areas may have low outbound demand at 4h but could serve as destinations for incoming trips. The placement of charging infrastructure without initial vehicle allocation in these areas demonstrates foresight in system design, preparing for potential incoming traffic and early morning demand.

5. Case Study




Figure 5.10.: Spatial Distribution at 4h

5. Case Study

Time Period 2 (9h) (see Figure 5.11)

The deployment plan for the 9h time period captures the morning peak in Manhattan's traffic. This time of day is marked by heavy commuter traffic as people head to work and school. Compared to the 4h period, there is a noticeable shift in how resources are allocated, showing how the system adapts to changing demand.

At 9h, the distribution of charging stations shows a slight reduction in coverage compared to 4h, with 17 out of 25 possible locations having stations, down from 19 at 4h. This more focused approach suggests an optimization strategy that concentrates resources in areas with the highest morning demand, like the Upper East Side South (Region 237) and Penn Station/Madison Sq West (Region 186), while still keeping good access across the city.

Vehicle allocation at this time remains at a total of 30 vehicles but with a notably different distribution reflecting the complex morning traffic patterns. For example, Upper East Side North (Region 236), Upper West Side South (Region 239), Lincoln Square East (Region 142), and Midtown North (Region 163): The number of vehicles allocated to these regions has increased compared to 4h. This larger allocation reflects the high volume of commuters leaving these mainly residential areas during the morning peak.

Several stations, like Clinton West (Region 50) and Lenox Hill East (Region 140), have charging stations but no vehicles allocated, similar to what was seen in some areas during the 4h period. This strategy likely reflects these areas' roles as destination points during the morning rush, where charging is needed for incoming vehicles but there's no need for outbound capacity. Additionally, compared to the 4h deployment plan, regions like Central Harlem (Region 41), East Harlem North (Region 74), and Financial District South (Region 88) have opted not to build charging stations. This could be due to the system prioritizing areas with higher demand, given the limited resources for deploying charging stations.

The 9h deployment plan effectively addresses the needs identified in our temporal demand analysis for the morning peak. It shows a clear shift towards supporting commuter patterns, with increased resources in residential areas for outbound traffic and in commercial areas for inbound traffic.

Time Period 3 (15h) (see Figure 5.12)

The deployment plan for the 15h time period reflects Manhattan's mid-afternoon traffic patterns, which are typically a mix of business travel, shopping, and leisure activities.

The distribution of charging stations at 15h remains the same as at 9h, with 17 out of 25 potential locations having stations. Similarly, the total vehicle allocation also stays at 30, though there are some changes in how the vehicles are distributed across locations.

- Penn Station/Madison Sq West (Region 186): Remains a critical hub with 4 vehicles allocated, a slight decrease from 5 in the morning peak. This adjustment likely reflects a slightly reduced but still significant demand at this major transportation center during mid-afternoon hours.
- Midtown Center (Region 161), Clinton East (Region 48), Upper East Side North (Region 236), and Upper East Side South (Region 237): Each allocated 3 vehicles. This consistent allocation across these diverse areas (commercial, mixed-use, and residential) suggests a balanced demand during this time of day, possibly due to a mix of business activities, shopping, and leisure travel.
- Central Park (Region 43): Now has no vehicles allocated, down from 1 at 9h. This could indicate a shift in recreational patterns, with fewer people starting their park activities in the mid-afternoon.
- Clinton West (Region 50): Now has 1 vehicle allocated, whereas it had none at 9h. This suggests an increase in mid-day demand in this area, possibly due to business or leisure activities.

Interestingly, several stations, including Greenwich Village North (Region 113), and Two Bridges/Seward Park (Region 232), maintain their morning allocations of 1 vehicle each, indicating consistent demand throughout the day in these areas.

5. Case Study




Figure 5.12.: Spatial Distribution at 15h

5. Case Study

Time Period 4 (18h) (see Figure 5.13)

The deployment plan for the 18h time period reflects the evening peak in Manhattan's traffic patterns. This time is characterized by heavy commuter traffic as people return from work and head out for evening activities, resulting in another notable shift in resource allocation compared to the mid-afternoon period (15h).

Charging station distribution at 18h is 18 out of 25 potential locations having stations built. And vehicle allocation at this time remains consistent with the 15h period, but with notable adjustments to address the evening rush hour patterns:

- Midtown Center (Region 161): Allocated 4 vehicles, an increase from 3 at 15h. This boost likely reflects the high volume of commuters leaving the city's principal commercial district during the 18h rush hour.
- Clinton East (Region 48) and East Village (Region 79): Both allocated 3 vehicles each. For Clinton East, this maintains the 15h allocation, while for East Village, this represents an increase. This allocation likely addresses the mix of residents returning home and the influx of people for evening entertainment in these vibrant areas.
- Penn Station/Madison Sq West (Region 186): Decreased from 4 to 2 vehicles, possibly indicating a shift from incoming travelers at 15h to outbound commuters at 18h.

The 18h deployment plan effectively addresses the needs identified in our temporal demand analysis for the evening peak. It shows a clear shift towards supporting reverse commute patterns, with increased resources in commercial areas for outbound traffic and adjustments in residential and entertainment areas to accommodate returning residents and evening activities.

5. Case Study

Time Period 5 (21h) (see Figure 5.14)

The deployment plan for the 21h time period reflects the late evening characteristics of Manhattan's traffic patterns. This period is typically characterized by a mix of late-night entertainment, dining activities, and the final wave of returning commuters.

Charging station distribution at 21h remains consistent with the 18h period, with 18 out of 25 potential locations having stations built. This consistency in the number of stations suggests that the spatial coverage established for the evening peak (18h) remains suitable for addressing late-night mobility needs. However, there are some notable changes in the specific locations of these stations, indicating a fine-tuning of the system to meet shifting demand patterns. Vehicle allocation at 21h maintains a total of 30 vehicles but with notable adjustments to address the late evening patterns:

- Midtown Center (Region 161): Maintains the highest allocation of 4 vehicles, unchanged from 18h. This consistency indicates the area's sustained importance as a hub for late-night activities and potentially late-working professionals.
- Penn Station/Madison Sq West (Region 186): Increased to 3 vehicles from 2 at 18h, possibly reflecting an uptick in late-night arrivals and departures at this major transportation hub.
- East Village (Region 79): Decreased from 3 to 1 vehicle, which is surprising given its reputation for nightlife. This might suggest a shift in late-night travel patterns or a focus on incoming rather than outgoing trips in this area.
- Central Harlem (Region 41): no longer has a charging station built, a change from the 18h period.

The 21h deployment plan strategically addresses the demands of the late evening hours, balancing the needs for entertainment, dining, and late-night commuting. It reflects a thoughtful approach to resource allocation, adjusting support for nightlife areas while ensuring adequate coverage in residential zones for individuals returning home late.

5. Case Study




Figure 5.14.: Spatial Distribution at 21h

Overall Evaluation

As shown in table 5.7, we assess the effectiveness of the charging station deployment plan by analyzing key metrics across different time periods.

Table 5.7.: Summary of Station Metrics at Different Times

Time (h)	Station Number	Cars Number	Profit	ROI	DSR	CSU
4	19	29	1826.29	24.40	0.7487	0.1317
9	17	30	6688.99	97.88	0.1279	0.0240
15	17	30	6849.31	100.22	0.1153	0.0226
18	18	30	7070.68	98.53	0.1086	0.0233
21	18	30	6976.14	97.21	0.1110	0.0212

Economic performance across time periods (Profit and ROI)

The economic performance of the deployment plan varies significantly across different time periods, as evidenced by the profit and Return on Investment (ROI) metrics.

At 4h, the deployment plan shows the lowest profit (1826.29) and ROI (24.40). During this period, the demand level is relatively low, and the income of car-sharing service significantly exceeds the cost, therefore, the car-sharing system tends to deploy more charging stations and vehicle resources to serve more potential requested trips. Consequently, the value of profit and ROI becomes more sensitive to traffic demand levels. In contrast, the deployment plan at peak hours, particularly at 18h (evening peak), shows the highest profit (7070.68) and a high ROI (98.53). This demonstrates that the system's resource allocation is well-aligned with the high demand during evening rush hours, maximizing the income from car-sharing services relative to the costs. It is important to highlight that, although the 18h period generates the highest profit, the 15h period exhibits the highest return on investment (ROI) at 100.22. This indicates that the afternoon timeframe may represent an optimal balance between demand and resource allocation, allowing the system to achieve peak efficiency in terms of ROI.

Customer satisfaction and infrastructure utilization (DSR and CSU)

The DSR and CSU offer valuable insights into the system's service efficiency and resource utilization across various time periods. At 4h, both DSR (0.7487) and CSU (0.1317) reach their peak values, significantly higher than at other time periods. This high efficiency during low-demand hours is a direct result of the relationship between resource allocation and demand level. Despite the low demand, the system maintains a high number of charging stations (19) and vehicles (29), nearly equivalent to peak hour allocations. In contrast, during other time periods (9h, 15h, 18h, 21h), we observe significantly lower DSR (ranging from 0.1086 to 0.1279) and CSU (ranging from 0.0212 to 0.0240) values. This drop in both metrics during higher demand periods suggests that: The system struggles to meet the increased demand despite maintaining high resource levels.

5.4.4. Deployment Plan Recommendations

Table 5.8 presents our recommendations for the deployment plan of electric charging stations. For each region, we provide construction priority recommendations along with the

5. Case Study

corresponding vehicle resources. Specifically, for the five-time periods (4h, 9h, 15h, 18h, and 21h), regions where no charging station is proposed are marked as "Not Recommended."

In cases where charging stations are required only during specific time periods, we advise "Evaluate with Caution." For these locations, our recommendations are as follows:

- Conduct a more detailed demand analysis.
- Consider small-scale charging stations.
- Treat these areas as potential expansion sites for future system growth.

For locations requiring charging stations across all time periods, we assign a "High Priority" recommendation for construction. Additionally, we provide recommendations for the corresponding vehicle assignment levels at different times of the day to ensure optimal resource allocation and efficiency.

Table 5.8.: Deployment Plan Recommendations Table

Latitude	Longitude	Station	Region	Construction Priority	Vehicle Resource Recommendations
40.70454	-74.0143	12	Battery Park	Not Recommended	
40.80759	-73.9550	41	Central Harlem	Evaluate with Caution	
40.78152	-73.9627	43	Central Park	Evaluate with Caution	
40.76475	-73.9880	48	Clinton East	High Priority	4 vehicles at early morning (4h), 2-3 vehicles rest of day
40.77123	-73.9936	50	Clinton West	Evaluate with Caution	
40.74778	-74.0004	68	East Chelsea	High Priority	1 vehicle base level, 2 vehicles at 15h
40.80182	-73.9393	74	East Harlem North	Evaluate with Caution	
40.72605	-73.9835	79	East Village	High Priority	2 vehicles at 4h/9h, 3 vehicles at 18h, 1 vehicle other times
40.70347	-74.0116	88	Financial District South	Evaluate with Caution	
40.73468	-73.9946	113	Greenwich Village North	High Priority	Consistent 1 vehicle across all time periods
40.76118	-73.9579	140	Lenox Hill East	Evaluate with Caution	
40.77384	-73.9821	142	Lincoln Square East	High Priority	1 vehicle at 4h, 2 vehicles rest of day
40.75854	-73.9772	161	Midtown Center	High Priority	4 vehicles at 4h/18h/21h, 3 vehicles at 9h/15h
40.76630	-73.9819	163	Midtown North	High Priority	1 vehicle off-peak (4h/15h), 2 vehicles peak hours (9h/18h/21h)
40.74825	-73.9763	170	Murray Hill	High Priority	2 vehicles at 4h, 1 vehicle rest of day
40.74910	-73.9920	186	Penn Station/Madison Sq West	High Priority	5 vehicles at 4h/9h, 4 vehicles at 15h, 2-3 vehicles evening
40.78419	-73.9266	194	Randalls Island	Not Recommended	
40.71978	-74.0068	231	TriBeCa/Civic Center	Evaluate with Caution	
40.71524	-73.9842	232	Two Bridges/Seward Park	Evaluate with Caution	
40.73792	-73.9922	234	UN/Turtle Bay South	Evaluate with Caution	
40.78017	-73.9551	236	Upper East Side North	High Priority	2 vehicles at 4h, 3 vehicles rest of day
40.76413	-73.9688	237	Upper East Side South	High Priority	2 vehicles base level, 3 vehicles at 15h, 1 vehicle at 18h
40.78184	-73.9793	239	Upper West Side South	High Priority	3 vehicles at 9h, 2 vehicles midday, 1 vehicle early/late
40.85650	-73.9328	243	Washington Heights North	Not Recommended	
40.83697	-73.9401	244	Washington Heights South	Not Recommended	



6

Discussion

6. Discussion

This chapter provides a comprehensive discussion of the LMSP Framework's performance and its broader implications. The discussion is organized into three main sections: First, we analyze the effectiveness of the framework by conducting a detailed evaluation of the results and deployment strategies (see section 6.1). Second, we critically examine the assumptions and limitations that underlie our LMSP Framework (see section 6.2). Finally, we assess the framework's generalizability across different urban contexts (see section 6.3). Through this systematic exploration, we aim to offer a balanced perspective on both the framework's contributions and potential areas for future improvement.

6.1. Discussion of Results

6.1.1. Effectiveness of the LMSP Framework

As demonstrated in the comparative analysis in section 5.3.5, the integration of deep learning and stochastic programming in our LMSP Framework significantly enhances the optimization of charging station locations. The innovative incorporation of deep learning methods allows for more precise demand prediction forecasts, which is crucial for addressing the inherent uncertainties in traffic patterns.

Our analysis shows that while the LMSP Framework may not yield the highest short-term profits compared to some traditional methods, it significantly outperforms these methods in terms of customer satisfaction, as measured by the Demand Satisfaction Ratio (DSR), and infrastructure efficiency, as indicated by Charging Station Utilization (CSU). Specifically, the LMSP Framework achieves an average DSR improvement of 26.32% and CSU improvement of 36.17% compared to the Historical Frequency-based Approach (HFA) across five time periods. Similarly, when compared to the Gaussian Distribution Approach (GDA), the LMSP Framework provides a 14.55% increase in DSR and a 12.06% increase in CSU.

These metrics indicate that our approach is better aligned with the actual needs of the system, particularly in environments characterized by high uncertainty and fluctuating demand patterns. While traditional methods may be more appropriate for scenarios focused on short-term profitability, the LMSP Framework provides a more balanced and sustainable approach. Its ability to adapt to the complex dynamics of urban traffic facilitates more effective resource allocation, resulting in higher levels of customer satisfaction and improved long-term operational efficiency for car-sharing systems.

6.1.2. Deployment Strategy Analysis

Section 5.4 presents optimal deployment plans for Manhattan from five specific time periods on January 1, 2020. These representative traffic patterns, to some extent, capture the daily cyclical nature of Manhattan's traffic conditions. However, it's important to acknowledge the limitations of this approach. The study does not account for operational activities, and the strategies provided are not real-time but rather static deployments under various specific time periods.

In our overall evaluation of these deployment plans across the five time periods (Table 5.7), an interesting observation emerges. Despite significant variations in demand levels, the system tends to allocate charging resources close to the maximum limit even during low demand hours, such as 4h, similar to high demand periods. This tendency can be attributed to the model's parameter settings, where the income per service significantly exceeds the cost parameters.

This finding highlights a crucial aspect of model design in transportation systems: the balance between resource allocation and demand levels. While ensuring sufficient resources during peak hours is essential, the over-allocation during off-peak hours suggests a potential for optimization. It indicates that the model might be overly aggressive in resource deployment, possibly leading to inefficiencies in low-demand periods.

This observation highlights the importance of fine-tuning model parameters to more accurately reflect real-world economic considerations. In practical applications, the cost-benefit ratio of deploying charging infrastructure should be carefully calibrated to avoid over-investment during low-utilization periods. This could involve introducing more refined cost structures or implementing dynamic pricing strategies that better align with the fluctuating demand levels throughout the day.

Additionally, this result highlights the need for more advanced modeling of operational costs and benefits. Future iterations of the model could integrate more detailed operational factors, such as the costs associated with maintaining underutilized stations and the potential advantages of having reserve capacity to accommodate unexpected demand surges.

In conclusion, while our LMSP Framework demonstrates significant advantages in optimizing charging station locations under demand uncertainty, the deployment strategy analysis reveals areas for further refinement. By addressing the issue of resource allocation across varying demand levels, future research can enhance the model's practical applicability and economic efficiency in real-world urban transportation systems.

6.2. Discussion of Assumptions and Limitations

6.2.1. Assumptions

While our proposed LMSP Framework provides valuable insights into optimizing charging station locations for electric car-sharing systems, it is important to acknowledge and discuss the assumptions made in our model and their potential limitations. This section examines these assumptions, their implications, and potential areas for future research.

Exclusion of Operational Activities

To simplify the modeling process of electric vehicle (EV) operations in a one-way car-sharing system, we made the assumption to exclude operational activities such as vehicle relocation and charging station scheduling costs. Instead, we employed a two-stage stochastic programming approach, which allowed us to focus on the strategic deployment of charging stations.

This assumption provided certain advantages, such as simplifying the model and enabling a concentrated focus on key infrastructure decisions, while also reducing computational complexity, thus making it feasible to solve larger-scale problems. However, this assumption also introduced some limitations. By excluding operational activities, the model may lack realism in representing the day-to-day operations of the system, which could result in suboptimal solutions regarding overall system efficiency. Additionally, the model's inability to capture short-term operational adjustments may hinder its ability to influence long-term strategic decisions effectively. Future research could address these limitations by incorporating operational activities using multi-stage stochastic programming or simulation studies, providing a more comprehensive view of system performance and leading to more robust optimization strategies.

Vehicle and Charging Station Uniformity

Our model assumes all vehicles are of the same type and charging stations have a fixed maximum capacity. While this assumption helps reduce model complexity, it doesn't reflect the diversity present in real-world scenarios. This simplification fails to capture the heterogeneity of EV fleets in real-world scenarios, ignores the potential benefits of dynamic capacity management at charging stations, and may lead to possible overestimation or underestimation of system capacity and efficiency. Future research could incorporate heterogeneous vehicle fleets and flexible charging station capacities to enhance model realism and capture potential operational efficiencies.

Simplified Origin-Destination Demand

In our model, we assume that all trips start and end at charging stations, ignoring the "first mile" and "last mile" of user journeys. In reality, users typically walk to and from charging stations at the beginning and end of their trips. This simplification potentially overestimates system convenience for users, may lead to suboptimal station placements in terms of user accessibility, and provides an inaccurate representation of total trip times and user experience. Future work could incorporate these walking distances into the demand modeling process, providing a more accurate representation of user behavior and system efficiency.

Simplified User Behavior Assumptions

Our model assumes users always choose the shortest driving path and park at charging stations. This simplification, while useful for modeling purposes, ignores the diversity of real user behaviors. It may lead to inaccurate representation of actual user route choices, ignore

the possibility of parking at non-charging locations, and potentially overestimate charging station utilization. Incorporating more diverse user behavior models, possibly through agent-based simulations, could enhance the realism of the system model and provide more accurate predictions of system performance.

Simplified Traffic Conditions

The model doesn't account for geographical factors or road conditions (such as congestion levels or road quality) that might affect traffic demand. It also doesn't consider potential queuing issues when returning vehicles. This simplification may result in potentially unrealistic travel time estimates, possible suboptimal station placements in areas with frequent congestion, underestimation of waiting times during peak hours, and an inability to capture the impact of road network characteristics on system performance. Future research could integrate more detailed traffic models and consider queuing theory to address these limitations. This could lead to more realistic travel time estimates and better-informed decisions on charging station placements.

6.2.2. Limitations

Although this research provides valuable insights into the optimization of charging station locations for electric car-sharing systems under demand uncertainty, it is still subject to certain limitations.

Firstly, this research relies heavily on historical data for predicting traffic demand. While this data provides a basis for forecasting it, it inherently lacks real-time updates and may not reflect sudden changes in traffic patterns due to unforeseen events or shifts in user behavior. Furthermore, the model does not incorporate geographic variables such as local geography, road conditions, or population density variations, all of which can significantly influence traffic flows and car-sharing demand. The exclusion of these factors can lead to inaccuracies in demand prediction.

Additionally, the research did not incorporate potential changes in regulatory policies or economic fluctuations that could significantly impact the demand for car-sharing services. Such external factors are critical in real-world applications and could affect the practical utility of the findings.

The use of deep learning models, while innovative, brings its own set of limitations. These models require large datasets to train effectively and are often seen as "black boxes," offering little in terms of the interpretability of the factors driving the predictions. The model's reliance solely on historical data without considering real-time data feeds or integrating multidisciplinary factors such as urban planning and socio-economic data may reduce the robustness of the predictive outcomes.

The stochastic programming model used in this research sampled only five scenarios to represent the variability in traffic demand, which may not sufficiently capture the full spectrum of potential traffic conditions. This limited sampling could undermine the model's ability to generalize and handle real-world uncertainties effectively. As such, the outcomes might not reflect the true variability and complexity of demand patterns that would be encountered in practice.

Furthermore, the probabilistic models used for stochastic programming are based on simplifying assumptions necessary for computational feasibility. These assumptions might not fully capture the complex inter-dependencies and variability in urban traffic patterns, potentially limiting the model's accuracy and real-world applicability. This can result in recommendations that are theoretically optimal but may not be as effective when implemented in varying real-world conditions.

Lastly, while the predictive deep learning model effectively forecasts hourly traffic demand, this level of granularity may not fully align with the requirements of long-term, strategic planning. The model provides several optimal deployment plans based on traffic demand at five specific hours in one day, and although our analysis offers conclusions that consider daily traffic fluctuations, our deployment recommendations are insufficient in reflecting seasonal shifts, holiday periods, or other long-term patterns that influence transportation behaviors over seasons or years. This limitation underscores the challenge of applying short-term predictive insights to long-term infrastructure decisions, which may lead to sub-optimal strategic outcomes.

6.3. Generalisability of the model and results

The generalisability of our LMSP Framework is an important consideration for its potential application in various urban environments. This section discusses the model's adaptability to different contexts and the factors that may influence its effectiveness in other settings.

6.3.1. Applicability to Urban Areas

Our LMSP Framework, based on research conducted in Manhattan, demonstrates excellent performance in high-density urban environments. It is particularly well-suited for cities with similar characteristics, such as high traffic demand and significant daily fluctuations in travel patterns. The model's effectiveness in capturing and responding to these dynamic urban conditions makes it a valuable tool for cities facing similar transportation challenges. However, its applicability to low-density areas or smaller cities requires further investigation. The unique traffic patterns, infrastructure, and user behaviors in less densely populated areas may necessitate adjustments to the model's parameters and assumptions.

6.3.2. Scalability to Other Cities

The flexibility of our model, particularly the deep learning component (LSTM-MLP-MDN), allows for potential application across multiple datasets for traffic demand prediction in different cities. This adaptability is a significant strength of the framework. However, it's important to note that the model may require retraining and fine-tuning to account for local factors specific to each city. These factors could include different regulations, geographical challenges, or unique user behaviors that influence traffic patterns and car-sharing system usage. The process of adapting the model to new urban environments would involve not only retraining with local data but also potentially adjusting the model architecture to capture city-specific features effectively.

6.3.3. Data Availability and Model Training

Our model relies heavily on historical traffic data, with the NYC taxi dataset serving as an excellent training set due to its comprehensiveness. The effectiveness of the model in new locations is contingent on the availability of similar high-quality, comprehensive datasets. In areas lacking such extensive data, the predictive accuracy of the model may be compromised. This limitation underscores the importance of data infrastructure and collection in cities looking to implement such advanced transportation planning tools. Future research could explore methods to adapt the model to work with less comprehensive datasets or to incorporate alternative data sources that might be more readily available in different urban contexts.

6.3.4. Limitations of Generalisation to Short-term or Real-time Adjustments

While our model excels in strategic planning, it currently does not account for short-term operational adjustments such as real-time vehicle repositioning or changes in traffic patterns due to unpredictable events like concerts or sports events. This limitation may affect its applicability in scenarios where rapid, dynamic responses to changing conditions are crucial. Extending the model to incorporate real-time adjustments could significantly enhance its applicability across various domains and improve its responsiveness to the day-to-day fluctuations in urban transportation needs.

In conclusion, while our LMSP Framework shows promising potential for generalization, particularly in high-density urban environments, its application in different contexts requires careful consideration of local factors, data availability, and the need for short-term operational flexibility. Future research directions could focus on enhancing the model's adaptability to diverse urban settings and incorporating real-time adjustment capabilities to broaden its applicability and effectiveness across a wider range of urban transportation scenarios.



7

Conclusion

7. Conclusion

7.1. Answers to the Research Questions

This section will address the questions presented in section 1.3.

7.1.1. Main Research Question

How can a novel deep learning-based stochastic programming framework (LMSP Framework) be developed and applied to optimize the Charging Station Location Problem (CSLP) in a one-way electric car-sharing system under conditions of traffic demand uncertainty?

Answer: This research successfully develops and implements a novel approach that combines deep learning with stochastic programming (LMSP Framework) to optimize the Charging Station Location Problem (CSLP) in one-way electric car-sharing systems under demand uncertainty. The key components and steps of this framework are as follows:

- **Development of Deep Learning Models:** Advanced deep learning models, including Long Short-Term Memory (LSTM), Multilayer Perceptron (MLP), and Mixture Density Networks (MDN), were used to predict traffic demand patterns more accurately.
- **Integration with Stochastic Programming:** These traffic demand predictions were integrated into a two-stage stochastic programming model, allowing for strategic decisions about the location and capacity of charging stations under different demand scenarios.
- **Optimization Under Uncertainty:** By considering demand uncertainty, the framework optimized the locations of charging stations and the allocation of vehicles. This method made the planning process more adaptable and robust in the face of unpredictable traffic.
- **Case Study Application:** The framework was tested with a detailed case study in Manhattan, demonstrating how it could be applied effectively in a real-world, high-density urban environment.
- **Performance Evaluation:** The framework was evaluated based on various criteria, such as profitability, return on investment, demand satisfaction ratio, and charging station utilization. These metrics showed that the framework was able to strike a good balance between economic efficiency and operational performance.
- **Comparative Analysis:** When compared to traditional approaches, the LMSP Framework showed better performance, particularly in improving customer satisfaction and making more efficient use of infrastructure, especially in areas where demand is highly variable and uncertain.

7.1.2. Sub Research Questions

1. How does the integration of stochastic programming help in accommodating demand uncertainty when optimizing the location and initial number of cars at charging stations?

Answer: The integration of stochastic programming in our LMSP Framework plays a crucial role in accommodating demand uncertainty for charging station optimization. By employing a two-stage stochastic programming model, our approach effectively handles the variability in traffic demand patterns. The first stage makes strategic decisions on charging station locations and initial vehicle allocations, while the second stage addresses operational decisions based on realized demand scenarios. This structure allows the model to consider multiple possible demand outcomes, each with associated probabilities when making infrastructure decisions. Consequently, the resulting charging station network is more robust and adaptable to fluctuating demand patterns. This approach significantly improves upon deterministic models by minimizing the risk of over- or under-provisioning resources, thereby enhancing the overall system's efficiency and resilience to demand uncertainty.

2. How can advanced deep learning techniques, particularly the LSTM-MLP-MDN model, be employed to accurately forecast the probability distribution of traffic demand in a one-way electric car-sharing system?

Answer: The research demonstrates that the integrating of advanced deep learning techniques, specifically the LSTM-MLP-MDN models (see Figure 4.2), significantly enhances the accuracy of predicting traffic demand probability distribution in one-way electric car-sharing system. As evidenced in table 5.3, the model achieves superior performance with an MAE of 174.3023, MAPE of 15.10%, and R^2 of 0.8872, significantly outperforming traditional forecasting methods.

The LSTM component of the model effectively captures the temporal dependencies in traffic demand. By processing sequences of historical demand data, LSTM can predict future demand patterns with greater accuracy, particularly by recognizing demand trends and fluctuations over time. This capability is essential for managing the dynamic nature of demand in a car-sharing system. The MLP adds a layer of complexity by modeling the nonlinear relationship between various influencing factors, such as weather, time of day, and traffic conditions, and their impact on demand. By doing so, the MLP refines the predictions generated by the LSTM, ensuring that the models account for complex patterns and interactions within the data. The MDN component is crucial for addressing the inherent uncertainty in traffic demand. Unlike traditional models that provide a single-point estimate, the MDN generates a probability distribution of potential demand levels. This probabilistic approach allows the model to not only predict the most likely demand scenario but also to account for a range of possible outcomes, thereby improving the robustness of the decision-making process in station location and vehicle allocation.

3. How does the integration of advanced demand prediction techniques, such as LSTM-MLP-MDN, enhance the performance of stochastic programming models addressing the CSLP compared to traditional forecasting methods, and what added value does it bring?

The integration of our advanced demand prediction model significantly enhances the optimization performance of stochastic programming models in solving the CSLP. Specifically, compared to the Historical Frequency-based Approach (HFA), the LMSP Framework achieves an average improvement of 26.32% in Demand Satisfaction Ratio (DSR) and 36.17% in Charging Station utilization rate (CSU) across five time periods. When compared to the

7. Conclusion

Gaussian Distribution Approach (GDA), it shows an average increase of 14.55% in DSR and 12.06% in CSU. The LSTM-MLP-MDN model's ability to generate more accurate probability distributions of traffic demand provides the stochastic programming component with higher quality input data, resulting in more robust and reliable optimization outcomes.

While traditional methods show higher short-term profits and return on investment due to their aggressive demand predictions, the LMSP Framework's superior performance in operational metrics indicates its ability to allocate resources more rationally and avoid waste. This demonstrates that the integration of advanced demand prediction techniques with stochastic programming brings significant added value to the planning and operation of electric car-sharing systems. The framework's focus on operational efficiency and demand satisfaction, rather than just short-term financial gains, positions it as a more viable option for long-term profitability and sustainable system operation, particularly in addressing the complex challenges of demand uncertainty in urban transportation systems.

4. What insights can be drawn from the application of the proposed LMSP Framework in a real-world case study (Manhattan)?

Answer: The application of the LMSP Framework in Manhattan provides several valuable insights into the optimization of charging station locations for electric car-sharing systems. Through comprehensive temporal and spatial analysis of traffic patterns, the following key findings emerge:

- The research reveals distinct traffic demand patterns throughout the day. Significant variations were observed across different time periods, such as 4h, 9h peak hour, 15h, 18h evening peak, and 21h. This time sensitivity highlights the importance of considering temporal factors in urban car-sharing system planning.
- Spatial analysis uncovers significant geographical variations in demand patterns across Manhattan. High-demand regions, such as Upper East Side South (Region 237) and Midtown Center (Region 161), consistently show substantial traffic demand, while areas like Randalls Island (Region 194) and Two Bridges/Seward Park (Region 232) exhibit consistently lower demand. These spatial variations emphasize the importance of location-specific deployment strategies. Furthermore, different urban characteristics significantly influence traffic patterns. Residential areas like Upper East Side South show strong morning and evening peaks, while commercial districts such as Midtown Center experience high daytime demand. Mixed-use areas like Clinton East demonstrate more consistent demand throughout the day.
- The deployment plan analysis reveals important insights about resource allocation. Our study demonstrates that optimal charging station placement and vehicle allocation vary significantly based on both location and time of day. The system's performance, measured through Demand Satisfaction Ratio (DSR) and Charging Station Utilization (CSU) rates, shows notable variations across different periods, reflecting the dynamic nature of urban mobility demands. While traditional methods show higher short-term profits, the LMSP Framework demonstrates superior operational efficiency, achieving better resource utilization and demand satisfaction. This suggests improved long-term sustainability and operational viability.
- These insights from the Manhattan case study not only validate the effectiveness of the LMSP Framework but also provide valuable guidance for urban planners and transportation system operators. The findings highlight the importance of considering both

7. Conclusion

temporal and spatial factors in charging infrastructure planning, while also demonstrating the potential of data-driven approaches in addressing complex urban mobility challenges. This comprehensive understanding of demand patterns and system performance contributes significantly to the development of more efficient and sustainable urban transportation solutions.

These findings not only validate the performance of the LMSP framework but also provide valuable insights for urban planners, policymakers, and transportation researchers. They highlight the complexity of urban mobility systems and the potential for data-driven, AI-enhanced tools to address these challenges effectively.

5. How do the results of the case study validate the effectiveness, adaptability, and feasibility of the proposed LMSP Framework?

Answer: The application of the LMSP Framework in the Manhattan case study reveals several important insights and effectively validates its effectiveness, adaptability, and feasibility in a real-world urban environment.

- **Effectiveness:** The LMSP Framework's effectiveness is primarily demonstrated by its significant improvement in demand forecasting accuracy and resource optimization.
 - **Demand Prediction:** By integrating deep learning techniques (LSTM-MLP-MDN), the LMSP Framework outperforms traditional methods in traffic demand prediction, achieving lower Mean Absolute Error (MAE) and higher R^2 values. Specifically, the LMSP Framework achieved an MAE of 174.3023, MAPE of 15.10%, and R^2 of 0.8872, significantly outperforming other comparative models.
 - **Resource Optimization:** More accurate predictions enable the framework to make better decisions on charging station placements and vehicle allocations, resulting in higher Demand Satisfaction Ratios (DSR) and Charging Station utilization rates (CSU) across different time periods. For example, it performs well during both low-demand (4 AM) and high-demand (6 PM) periods.
- **Adaptability:** The LMSP Framework's adaptability is highlighted by its ability to handle diverse and fluctuating demand patterns. This adaptability confirms the framework's capability to respond to both temporal and spatial variations in traffic, crucial for managing demand uncertainty in complex urban environments.
 - **Temporal Adaptability:** The framework adapts well to both low-demand periods (e.g., 4 AM) and high-demand periods (e.g., 6 PM) throughout the day. It provides reasonable charging station layouts and vehicle allocations for five specific time periods (4h, 9h, 15h, 18h, 21h).
 - **Spatial Adaptability:** It successfully captures and adapts to demand differences between commercial areas (like Midtown Center) and residential areas (such as Upper East Side South).
- **Feasibility:** The feasibility of the LMSP Framework is validated through its successful application in a large-scale, high-density urban environment like Manhattan:
 - **Scalability:** The LMSP Framework successfully optimizes a network of 25 potential charging station locations across Manhattan, incorporating various demand scenarios derived from extensive historical data. The model parameters, including charging station costs, vehicle purchase costs, and operational metrics, were

7. Conclusion

calibrated using real-world data from New York City Taxi & Limousine Commission reports. This demonstrates the framework's ability to handle large-scale, data-intensive problems in a complex urban environment.

- Practical Application: The framework's recommendations for station placements in high-demand areas such as Midtown Center and Upper East Side South align with known traffic patterns, validating its practical applicability.
- Computational Efficiency: Solves large-scale optimization problems within reasonable time frames, demonstrating its feasibility in real-world scenarios.

5. How can the findings from this research enhance the planning and operation of electric car-sharing systems, and what are the potential implications for sustainable urban mobility?

This research significantly enhances the planning and operation of electric car-sharing systems through the LMSP Framework's improved optimization of charging station locations and vehicle allocations. The framework's ability to handle demand uncertainty, as evidenced by superior DSR and CSU metrics, can lead to more efficient and reliable services. Insights from temporal and spatial demand analysis can inform broader urban planning decisions, integrating car-sharing systems more effectively into overall transportation networks. The implications for sustainable urban mobility are substantial: the optimized systems can contribute to reducing private vehicle ownership, decreasing urban congestion, and lowering emissions. By promoting efficient resource utilization and aligning with sustainability goals, this research has the potential to accelerate the transition to cleaner transportation options in urban areas globally, contributing to more sustainable and livable cities.

7.2. Future Research Outlook

The findings of this study open up several avenues for future research in the field of electric car-sharing systems and charging station optimization. These potential research directions aim to enhance the applicability, accuracy, and comprehensiveness of the LMSP Framework.

1. Geographical Expansion: Future research should expand the geographical scope of case studies beyond Manhattan to validate the proposed model in different urban and suburban environments. This would help assess the generalizability of the model's conclusions and adapt the optimization strategies to cities with varying transportation infrastructures and population densities. Such studies could provide valuable insights into how the model performs under diverse urban conditions and help refine its parameters for broader applicability.
2. Integration of Real-time Data: Integrating real-time traffic data and geographic variables such as road conditions and population density could significantly enhance the accuracy of demand forecasts. This integration would lead to more dynamic and adaptable car-sharing systems, capable of responding to immediate changes in traffic patterns or urban dynamics. Future studies could explore methods to incorporate these real-time data streams into the LMSP Framework, potentially through the development of online learning algorithms or adaptive optimization techniques.

7. Conclusion

3. **Operational Dynamics Integration:** Future research could focus on integrating short-term operational dynamics, such as vehicle relocation strategies and dynamic pricing models, into the long-term planning framework. This integration could provide a more holistic optimization approach, bridging the gap between strategic planning and day-to-day operations of car-sharing systems.
4. **External Factor Analysis:** Future work could explore the impact of external factors like regulatory changes and economic fluctuations on car-sharing demand. By incorporating scenario analyses that consider these variables, future research could provide more comprehensive insights into the long-term sustainability and adaptability of electric car-sharing services. This could involve developing more sophisticated economic models or policy simulation frameworks to be integrated with the existing LMSP Framework.



8

Appendices

A. Deployment Plan over Five Time Periods

A. Deployment Plan over Five Time Periods

Table A.1.: Deployment Plan at 4h

Region ID	y.i	L.i	Latitude	Longitude	Zone
12	0	0	40.7045416	-74.0142564	Battery Park
41	1	0	40.8075917	-73.9549547	Central Harlem
43	1	0	40.7815194	-73.9627249	Central Park
48	1	4	40.7647531	-73.9880232	Clinton East
50	1	0	40.7712271	-73.9935659	Clinton West
68	1	1	40.7477804	-74.000438	East Chelsea
74	1	0	40.80182	-73.9392953	East Harlem North
79	1	2	40.7260512	-73.9835308	East Village
88	1	1	40.7034746	-74.0115948	Financial District South
113	1	1	40.7346805	-73.9946012	Greenwich Village North
140	1	1	40.7611816	-73.9579018	Lenox Hill East
142	1	1	40.7738424	-73.9821234	Lincoln Square East
161	1	4	40.7585437	-73.9772064	Midtown Center
163	1	1	40.7663007	-73.9818927	Midtown North
170	1	2	40.7482478	-73.9762946	Murray Hill
186	1	5	40.749101	-73.992006	Penn Station/Madison Sq West
194	0	0	40.7841865	-73.9266152	Randalls Island
231	0	0	40.7197847	-74.0068153	TriBeCa/Civic Center
232	0	0	40.7152425	-73.9842337	Two Bridges/Seward Park
234	1	1	40.7379242	-73.9922478	UN/Turtle Bay South
236	1	2	40.7801748	-73.9550942	Upper East Side North
237	1	2	40.76413	-73.9688047	Upper East Side South
239	1	1	40.7818426	-73.979274	Upper West Side South
243	0	0	40.856503	-73.932761	Washington Heights North
244	0	0	40.8369673	-73.9401365	Washington Heights South

A. Deployment Plan over Five Time Periods

Table A.2.: Deployment Plan at 9h

Region ID	y.i	L.i	Latitude	Longitude	Zone
12	0	0	40.7045416	-74.0142564	Battery Park
41	0	0	40.8075917	-73.9549547	Central Harlem
43	1	1	40.7815194	-73.9627249	Central Park
48	1	2	40.7647531	-73.9880232	Clinton East
50	1	0	40.7712271	-73.9935659	Clinton West
68	1	1	40.7477804	-74.000438	East Chelsea
74	0	0	40.80182	-73.9392953	East Harlem North
79	1	2	40.7260512	-73.9835308	East Village
88	0	0	40.7034746	-74.0115948	Financial District South
113	1	1	40.7346805	-73.9946012	Greenwich Village North
140	1	0	40.7611816	-73.9579018	Lenox Hill East
142	1	2	40.7738424	-73.9821234	Lincoln Square East
161	1	3	40.7585437	-73.9772064	Midtown Center
163	1	2	40.7663007	-73.9818927	Midtown North
170	1	1	40.7482478	-73.9762946	Murray Hill
186	1	5	40.749101	-73.992006	Penn Station/Madison Sq West
194	0	0	40.7841865	-73.9266152	Randalls Island
231	0	0	40.7197847	-74.0068153	TriBeCa/Civic Center
232	1	1	40.7152425	-73.9842337	Two Bridges/Seward Park
234	1	1	40.7379242	-73.9922478	UN/Turtle Bay South
236	1	3	40.7801748	-73.9550942	Upper East Side North
237	1	2	40.76413	-73.9688047	Upper East Side South
239	1	3	40.7818426	-73.979274	Upper West Side South
243	0	0	40.856503	-73.932761	Washington Heights North
244	0	0	40.8369673	-73.9401365	Washington Heights South

A. Deployment Plan over Five Time Periods

Table A.3.: Deployment Plan at 15h

Region ID	y.i	L.i	Latitude	Longitude	Zone
12	0	0	40.7045416	-74.0142564	Battery Park
41	0	0	40.8075917	-73.9549547	Central Harlem
43	1	0	40.7815194	-73.9627249	Central Park
48	1	3	40.7647531	-73.9880232	Clinton East
50	1	1	40.7712271	-73.9935659	Clinton West
68	1	2	40.7477804	-74.000438	East Chelsea
74	0	0	40.80182	-73.9392953	East Harlem North
79	1	1	40.7260512	-73.9835308	East Village
88	0	0	40.7034746	-74.0115948	Financial District South
113	1	1	40.7346805	-73.9946012	Greenwich Village North
140	1	1	40.7611816	-73.9579018	Lenox Hill East
142	1	2	40.7738424	-73.9821234	Lincoln Square East
161	1	3	40.7585437	-73.9772064	Midtown Center
163	1	1	40.7663007	-73.9818927	Midtown North
170	1	1	40.7482478	-73.9762946	Murray Hill
186	1	4	40.749101	-73.992006	Penn Station/Madison Sq West
194	0	0	40.7841865	-73.9266152	Randalls Island
231	0	0	40.7197847	-74.0068153	TriBeCa/Civic Center
232	1	1	40.7152425	-73.9842337	Two Bridges/Seward Park
234	1	1	40.7379242	-73.9922478	UN/Turtle Bay South
236	1	3	40.7801748	-73.9550942	Upper East Side North
237	1	3	40.76413	-73.9688047	Upper East Side South
239	1	2	40.7818426	-73.979274	Upper West Side South
243	0	0	40.856503	-73.932761	Washington Heights North
244	0	0	40.8369673	-73.9401365	Washington Heights South

A. Deployment Plan over Five Time Periods

Table A.4.: Deployment Plan at 18h

Region ID	y.i	L.i	Latitude	Longitude	Zone
12	0	0	40.7045416	-74.0142564	Battery Park
41	1	0	40.8075917	-73.9549547	Central Harlem
43	1	1	40.7815194	-73.9627249	Central Park
48	1	3	40.7647531	-73.9880232	Clinton East
50	1	1	40.7712271	-73.9935659	Clinton West
68	1	1	40.7477804	-74.000438	East Chelsea
74	0	0	40.80182	-73.9392953	East Harlem North
79	1	3	40.7260512	-73.9835308	East Village
88	0	0	40.7034746	-74.0115948	Financial District South
113	1	1	40.7346805	-73.9946012	Greenwich Village North
140	1	1	40.7611816	-73.9579018	Lenox Hill East
142	1	2	40.7738424	-73.9821234	Lincoln Square East
161	1	4	40.7585437	-73.9772064	Midtown Center
163	1	2	40.7663007	-73.9818927	Midtown North
170	1	1	40.7482478	-73.9762946	Murray Hill
186	1	2	40.749101	-73.992006	Penn Station/Madison Sq West
194	0	0	40.7841865	-73.9266152	Randalls Island
231	0	0	40.7197847	-74.0068153	TriBeCa/Civic Center
232	1	1	40.7152425	-73.9842337	Two Bridges/Seward Park
234	1	1	40.7379242	-73.9922478	UN/Turtle Bay South
236	1	3	40.7801748	-73.9550942	Upper East Side North
237	1	1	40.76413	-73.9688047	Upper East Side South
239	1	2	40.7818426	-73.979274	Upper West Side South
243	0	0	40.856503	-73.932761	Washington Heights North
244	0	0	40.8369673	-73.9401365	Washington Heights South

A. Deployment Plan over Five Time Periods

Table A.5.: Deployment Plan at 21h

Region ID	y.i	L.i	Latitude	Longitude	Zone
12	0	0	40.7045416	-74.0142564	Battery Park
41	0	0	40.8075917	-73.9549547	Central Harlem
43	1	1	40.7815194	-73.9627249	Central Park
48	1	3	40.7647531	-73.9880232	Clinton East
50	1	1	40.7712271	-73.9935659	Clinton West
68	1	1	40.7477804	-74.000438	East Chelsea
74	0	0	40.8018200	-73.9392953	East Harlem North
79	1	1	40.7260512	-73.9835308	East Village
88	0	0	40.7034746	-74.0115948	Financial District South
113	1	1	40.7346805	-73.9946012	Greenwich Village North
140	1	1	40.7611816	-73.9579018	Lenox Hill East
142	1	2	40.7738424	-73.9821234	Lincoln Square East
161	1	4	40.7585437	-73.9772064	Midtown Center
163	1	2	40.7663007	-73.9818927	Midtown North
170	1	1	40.7482478	-73.9762946	Murray Hill
186	1	3	40.7491010	-73.9920060	Penn Station/Madison Sq West
194	0	0	40.7841865	-73.9266152	Randalls Island
231	1	1	40.7197847	-74.0068153	TriBeCa/Civic Center
232	1	1	40.7152425	-73.9842337	Two Bridges/Seward Park
234	1	1	40.7379242	-73.9922478	UN/Turtle Bay South
236	1	3	40.7801748	-73.9550942	Upper East Side North
237	1	2	40.7641300	-73.9688047	Upper East Side South
239	1	1	40.7818426	-73.9792740	Upper West Side South
243	0	0	40.8565030	-73.9327610	Washington Heights North
244	0	0	40.8369673	-73.9401365	Washington Heights South

B. Trip Data at Different Locations over five specific time periods

B. Trip Data at Different Locations over five specific time periods

Table B.1.: Trip Data at Different Locations

Location ID	Zone	4h	9h	15h	18h	Total Trips
12	Battery Park	3	35	152	28	236
41	Central Harlem	236	1601	1337	1561	5857
43	Central Park	118	3741	6388	4567	16977
48	Clinton East	1866	7310	5546	8047	30917
50	Clinton West	585	2761	2238	2700	9936
68	East Chelsea	1422	4501	4984	6071	22153
74	East Harlem North	201	1635	1416	1794	5968
79	East Village	1791	3609	3212	4047	18638
88	Financial District	55	351	690	948	2704
113	Greenwich Village	314	3466	3437	4323	15541
140	Lenox Hill East	197	5203	4694	4484	16590
142	Lincoln Square East	328	6889	8799	10308	35415
161	Midtown Center	351	6236	11687	12687	41322
163	Midtown North	569	4910	8292	10593	31242
170	Murray Hill	539	7588	6766	8236	29103
186	Penn Station/Madison Sq	930	8660	7313	8625	33177
194	Randalls Island	0	5	10	5	27
231	TriBeCa/Civic Center	440	2261	2672	3310	11587
232	Two Bridges/Seward Park	120	228	196	275	1132
234	UN/Turtle Bay South	631	4602	8249	10388	32689
236	Upper East Side North	298	10259	12484	10941	38984
237	Upper East Side South	317	10397	13973	13296	45027
239	Upper West Side South	385	6298	7741	7459	26439
243	Washington Heights North	48	208	148	125	615
244	Washington Heights South	82	519	577	567	2093

C. Fundamentals of Deep Learning Model

C.1. LSTM unit

A comprehensive breakdown of the LSTM unit's structure is provided here.

Forget Gate:

The forget gate determines which information from the previous cell state should be discarded or kept. Current time input x_t and previous hidden state h_{t-1} are fed to the gate, and outputs a number between 0 and 1 for each number in cell state C_{t-1} . A 1 means "entirely keep this" while 0 means "entirely discard this".

$$f_t = \sigma \left(W_f \cdot [h_{t-1}, x_t] + b_f \right) \quad (C.1)$$

where:

σ is the sigmoid activation function, W_f represents the weight matrix associated with the forget gate, $[h_{t-1}, x_t]$ denotes the concatenation of the current input and the previous hidden state.

Input Gate:

The input gate decides which new information to add to the cell state. This involves two parts: a sigmoid layer which decides which values to update, and a tanh layer which creates a vector of new candidate value, \tilde{C}_t that could be added to the state.

$$i_t = \sigma (W_i \cdot [h_{t-1}, x_t] + b_i) \quad (C.2)$$

$$\tilde{C}_t = \tanh (W_C \cdot [h_{t-1}, x_t] + b_C) \quad (C.3)$$

Update Cell State:

$$C_t = f_t * C_{t-1} + i_t * \tilde{C}_t \quad (C.4)$$

where C_{t-1} represents the old cell state, it is updated to the new cell state C_t . This is performed by multiplying the old cell state by the forget gate's output and adding the product of the input gates' output and the candidate values.

Output Gate:

C. Fundamentals of Deep Learning Model




Figure C.1.: Mixture Density Network structure [Petrov and Repin, 2020]

$$o_t = \sigma(W_o \cdot [h_{t-1}, x_t] + b_o) \quad (\text{C.5})$$

$$h_t = o_t * \tanh(C_t) \quad (\text{C.6})$$

The output gate of an LSTM, denoted by o_t , is computed using a sigmoid activation function applied to the linear combination of the current input x_t and the previous hidden state h_{t-1} , along with b_o . Then, the output gate activation o_t then modulates the output from the cell state C_t , which processes through a tanh function to normalize its values. The final output hidden state h_t can be obtained by element-wise multiplication of o_t with $\tanh(C_t)$, effectively filtering the cell state information to produce the output relevant for the next time step.

C.2. MDN

The Mixture Density Network consists of two parts: a Neural Network and a Mixture Model. A neural network adopts any effective structure to convert inputs \mathbf{x} into learned features, and a mixture model is a type of probabilistic model constructed with a weighted sum of simpler distributions [?].

Figure C.1 illustrates the structure of the Mixture Density Network. Given input \mathbf{x} , the MDN outputs weights π_i , mean values μ_i , and variances σ_i of the i -th Gaussian distribution. The probability density function of the i -th Gaussian distribution is given by:

$$f_i(y) = \frac{1}{\sqrt{2\pi\sigma_i^2}} \exp\left(-\frac{(y - \mu_i)^2}{2\sigma_i^2}\right) \quad (\text{C.7})$$

Weighted by mixing coefficients π_i , N Gaussian distributions form a mixture density conditional probability density $p(\mathbf{y} | \mathbf{x})$, which is denoted as:

$$f_{\text{mixture}}(x) = \sum_{i=1}^N \pi_i \cdot f_i(y) \quad (\text{C.8})$$




Figure C.2.: ReLU Activation Function




Figure C.3.: Softplus Activation Function

C.3. Activation Function

Rectified Linear Unit (ReLU)

Rectified Linear Unit (ReLU) is a type of activation function commonly used in neural networks, especially in the hidden layers of deep learning models. It introduces non-linearity by outputting the input directly if it is positive; otherwise, it outputs zero, as shown in Figure C.2. It is defined mathematically as:

$$\text{ReLU}(x) = \max(0, x) \quad (\text{C.9})$$

ReLU is computationally efficient and helps mitigate the vanishing gradient problem, making it well-suited for deep learning models.

Softplus

The Softplus activation function is a smooth, differentiable alternative to the ReLU (Rectified Linear Unit) function, as shown in Figure C.3. The Softplus function is defined mathematically as:

$$\text{Softplus}(x) = \log(1 + e^x) \quad (\text{C.10})$$

Softmax

The Softmax function transforms a vector of real numbers into a probability distribution, where each output value represents the probability of the corresponding class. The values output by Softmax sum to 1, making them directly interpretable as probabilities.

D. Results of Comparison between LMSP Framework and Traditional-based SP

Table D.1.: Performance at 4h

	Profit	ROI	DSR	CSU
HFA	2312.36	30.7577	0.6252	0.1205
GDA	3896.11	54.5521	0.3756	0.0621
LMSP	1826.29	24.4026	0.7487	0.1317

Table D.2.: Performance at 9h

	Profit	ROI	DSR	CSU
HFA	6819.04	100.2800	0.1190	0.0225
GDA	7177.60	105.0278	0.1068	0.0212
LMSP	6688.99	97.8781	0.1279	0.0240

Table D.3.: Performance at 15h

	Profit	ROI	DSR	CSU
HFA	7093.08	99.3151	0.0971	0.0188
GDA	7260.76	101.6629	0.0922	0.0181
LMSP	6849.31	100.2241	0.1153	0.0226

Table D.4.: Performance at 18h

	Profit	ROI	DSR	CSU
HFA	7331.66	102.1692	0.1005	0.0207
GDA	7457.97	103.9293	0.0964	0.0194
LMSP	7070.68	98.5323	0.1086	0.0233

Table D.5.: Performance at 21h

	Profit	ROI	DSR	CSU
HFA	7425.26	103.4736	0.0935	0.0190
GDA	7281.54	101.4708	0.0989	0.0186
LMSP	6976.14	97.2148	0.1110	0.0212

A monochromatic blue-tinted photograph of a city skyline behind a dense forest, with a gondola on a body of water in the foreground. The skyline includes a tall, ornate building on the left, a modern skyscraper in the center, and a large, classical-style building on the right. The foreground shows a gondola with several people on a body of water.

9

Bibliography

Bibliography

- Alghamdi, T., Elgazzar, K., Bayoumi, M., Sharaf, T., and Shah, S. (2019). Forecasting traffic congestion using arima modeling. *2019 15th International Wireless Communications Mobile Computing Conference (IWCMC)*, pages 1227–1232.
- Baptista, P., Melo, S., and Rolim, C. (2014). Energy, environmental and mobility impacts of car-sharing systems. empirical results from lisbon, portugal. *Procedia - Social and Behavioral Sciences*, 111:28–37.
- Baron, O., Milner, J. M., and Naseraldin, H. (2011). Facility location: A robust optimization approach. *Production and Operations Management*, 20:772 – 785.
- Bayram, I. S. and Bayhan, S. (2020). Location analysis of electric vehicle charging stations for maximum capacity and coverage. *2020 IEEE 14th International Conference on*.
- Boujelben, M. K. and Gicquel, C. (2019). Efficient solution approaches for locating electric vehicle fast charging stations under driving range uncertainty. *Comput. Oper. Res.*, 109:288–299.
- Brandstätter, G., Kahr, M., and Leitner, M. (2017). Determining optimal locations for charging stations of electric car-sharing systems under stochastic demand. *Transportation Research Part B: Methodological*, 104:17–35.
- Chen, K., Chen, F., Lai, B., Jin, Z., Liu, Y., Li, K., Wei, L., Wang, P., Tang, Y., Huang, J., and Hua, X. (2020). Dynamic spatio-temporal graph-based cnns for traffic flow prediction. *IEEE Access*, 8:185136–185145.
- Chen, Y.-W., Cheng, C.-Y., Li, S.-F., and Yu, C.-H. (2018). Location optimization for multiple types of charging stations for electric scooters. *Appl. Soft Comput.*, 67:519–528.
- Chrobok, R., Kaumann, O., Wahle, J., and Schreckenberg, M. (2004). Different methods of traffic forecast based on real data. *Eur. J. Oper. Res.*, 155:558–568.
- City of New York (2024). Transportation in nyc. <https://wherewelive.cityofnewyork.us/explore-data/access-to-opportunity/transportation/>. Accessed: 2024-08-01.
- Clewlow, R. R. (2016). Carsharing and sustainable travel behavior: Results from the san francisco bay area. *Transport Policy*, 51:158–164.
- Cocca, M., Giordano, D., Mellia, M., and Vassio, L. (2019). Free floating electric car sharing: A data driven approach for system design. *IEEE Transactions on Intelligent Transportation Systems*, 20:4691–4703.
- Deb, S., Tammi, K., zhi Gao, X., Kalita, K., Mahanta, P., and Cross, S. (2021). A robust two-stage planning model for the charging station placement problem considering road traffic uncertainty. *IEEE Transactions on Intelligent Transportation Systems*, 23:6571–6585.

Bibliography

- Department, N. Y. C. P. (2024). Motor vehicle collisions - crashes. Accessed: 2024-07-29.
- Deza, A., Huang, K., and Metel, M. R. (2018). Charging station optimization for balanced electric car sharing. *Discret. Appl. Math.*, 308:187–197.
- Diao, Z., Zhang, D., Wang, X., Xie, K., He, S., Lu, X., and Li, Y. (2019). A hybrid model for short-term traffic volume prediction in massive transportation systems. *IEEE Transactions on Intelligent Transportation Systems*, 20:935–946.
- Dimara, A., Triantafyllidis, D., Krinidis, S., Kitsikoudis, K., Ioannidis, D., Valkouma, E., Skarvelakis, S., Antipas, S., and Tzovaras, D. (2021). Mlp for spatio-temporal traffic volume forecasting. In *2021 IEEE International IOT, Electronics and Mechatronics Conference (IEMTRONICS)*, pages 1–7. IEEE.
- Du, B., Hu, X., Sun, L., Liu, J., Qiao, Y., and Lv, W. (2020). Traffic demand prediction based on dynamic transition convolutional neural network. *IEEE Transactions on Intelligent Transportation Systems*, 22:1237–1247.
- Duan, P., Mao, G., Liang, W., and gan Zhang, D. (2019). A unified spatio-temporal model for short-term traffic flow prediction. *IEEE Transactions on Intelligent Transportation Systems*, 20:3212–3223.
- Duan, Z., Yang, Y., Zhang, K., Ni, Y., and Bajgain, S. (2018). Improved deep hybrid networks for urban traffic flow prediction using trajectory data. *IEEE Access*, 6:31820–31827.
- Fan, W. (2014). Optimizing strategic allocation of vehicles for one-way car-sharing systems under demand uncertainty. *Journal of the Transportation Research Forum*, 53:7–20.
- Faridimehr, S., Venkatachalam, S., and Chinnam, R. (2017). A stochastic programming approach for electric vehicle charging network design. *IEEE Transactions on Intelligent Transportation Systems*, 20:1870–1882.
- Febbraro, A. D., Sacco, N., and Saeednia, M. (2019). One-way car-sharing profit maximization by means of user-based vehicle relocation. *IEEE Transactions on Intelligent Transportation Systems*, 20:628–641.
- Firnborn, J. and Müller, M. (2011). What will be the environmental effects of new free-floating car-sharing systems? the case of car2go in ulm. *Ecological Economics*, 70(8):1519 – 1528. Cited by: 387.
- Formentin, S., Bianchessi, A. G., and Savaresi, S. (2015). On the prediction of future vehicle locations in free-floating car sharing systems. *2015 IEEE Intelligent Vehicles Symposium (IV)*, pages 1006–1011.
- Gao, K., Li, D., Chen, L., Geng, J., Gui, F., Cheng, Y., and Gu, Y. (2020). Predicting traffic demand matrix by considering inter-flow correlations. *IEEE INFOCOM 2020 - IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, pages 165–170.
- Giménez-Gaydou, D. A., Ribeiro, A. S., Gutiérrez, J., and Antunes, A. P. (2016). Optimal location of battery electric vehicle charging stations in urban areas: A new approach. *International Journal of Sustainable Transportation*, 10(5):393–405.
- Glott-Richter, M. (2016). Reclaim street space! – exploit the european potential of car sharing. *Transportation research procedia*, 14:1296–1304.

Bibliography

- Guan, J., Wang, W., Li, W., and Zhou, S. (2018). A unified framework for predicting kpis of on-demand transport services. *IEEE Access*, 6:32005–32014.
- Guo, G. and Kang, M. (2022). Rebalancing and charging scheduling with price incentives for car sharing systems. *IEEE Transactions on Intelligent Transportation Systems*, 23:18592–18602.
- Guo, K., Hu, Y., Qian, Z., Sun, Y., Gao, J., and Yin, B. (2020). Dynamic graph convolution network for traffic forecasting based on latent network of laplace matrix estimation. *IEEE Transactions on Intelligent Transportation Systems*, 23:1009–1018.
- Guo, Z., Li, T., Peng, S., and Zhang, H. (2021). Environmental and economic consequences of the incentive policy on electric vehicle industry: A cge based study in china. *Resources Conservation and Recycling*, 169:105542.
- Hagberg, A. A., Schult, D. A., and Swart, P. J. (2008). Networkx. <https://networkx.github.io/>. Python package for the creation, manipulation, and study of the structure, dynamics, and functions of complex networks.
- Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9:1735–1780.
- Hua, Y., Zhao, D., Wang, X., and Li, X. (2019). Joint infrastructure planning and fleet management for one-way electric car sharing under time-varying uncertain demand. *Transportation Research Part B: Methodological*.
- International Energy Agency (2022). *Global EV Outlook 2022: Securing supplies for an electric future*. International Energy Agency, Paris.
- Jiao, Z., Ran, L., Chen, J., Meng, H., and Li, C. (2017). Data-driven approach to operation and location considering range anxiety of one-way electric vehicles sharing system. *Energy Procedia*, 105:2287–2294.
- Kadri, A., Perrouault, R., Boujelben, M. K., and Gicquel, C. (2020). A multi-stage stochastic integer programming approach for locating electric vehicle charging stations. *Comput. Oper. Res.*, 117:104888.
- Ke, H., Chai, S., and Cheng, R.-Z. (2019). Does car sharing help reduce the total number of vehicles? *Soft Computing*, pages 1–14.
- Kim, S., Rasouli, S., Timmermans, H., and Yang, D. (2021). A scenario-based stochastic programming approach for the public charging station location problem. *Transportmetrica B: Transport Dynamics*, 10:340 – 367.
- Kumar, S. V. and Vanajakshi, L. (2015). Short-term traffic flow prediction using seasonal arima model with limited input data. *European Transport Research Review*, 7.
- Li, C., Zhang, L., Ou, Z., Wang, Q., Zhou, D., and Ma, J. (2022). Robust model of electric vehicle charging station location considering renewable energy and storage equipment. *Energy*, 238:121713.
- Li, W., Cao, J., Guan, J., Zhou, S., Liang, G., So, W. K. Y., and Szczecinski, M. (2019). A general framework for unmet demand prediction in on-demand transport services. *IEEE Transactions on Intelligent Transportation Systems*, 20(8):2820–2830.

Bibliography

- Li, X., Ma, J., Cui, J., Ghiasi, A., and Zhou, F. (2016). Design framework of large-scale one-way electric vehicle sharing systems: A continuum approximation model. *Transportation Research Part B-methodological*, 88:21–45.
- Lyu, Z., ElSaid, A., Karns, J., Mkaouer, M., and Desell, T. (2021). An experimental study of weight initialization and lamarckian inheritance on neuroevolution. In *Applications of Evolutionary Computation: 24th International Conference, EvoApplications 2021, Held as Part of EvoStar 2021, Virtual Event, April 7–9, 2021, Proceedings*, page 584–600, Berlin, Heidelberg. Springer-Verlag.
- Ma, X., Dai, Z., He, Z., Ma, J., Wang, Y., and Wang, Y. (2017). Learning traffic as images: A deep convolutional neural network for large-scale transportation network speed prediction. *Sensors (Basel, Switzerland)*, 17.
- Meng, W. and Kai, L. (2011). Optimization of electric vehicle charging station location based on game theory. *Proceedings 2011 International Conference on Transportation, Mechanical, and Electrical Engineering (TMEE)*, pages 809–812.
- MirHassani, S. A., Khaleghi, A., and Hooshmand, F. (2020). Two-stage stochastic programming model to locate capacitated ev-charging stations in urban areas under demand uncertainty. *EURO Journal on Transportation and Logistics*, 9:100025.
- New York City Taxi & Limousine Commission (2022). Charged up! tlc’s roadmap to electrifying the for-hire transportation sector in new york city. Accessed: 2024-08-07.
- Nourinejad, M. and Roorda, M. J. (2015). Carsharing operations policies: a comparison between one-way and two-way systems. *Transportation*, 42:497–518.
- Open Meteo (2024). Historical weather api. Accessed: 2024-07-29.
- Pan, T., Sumalee, A., Zhong, R., and Indra-Payoong, N. (2013). Short-term traffic state prediction based on temporal-spatial correlation. *IEEE Transactions on Intelligent Transportation Systems*, 14:1242–1254.
- Petrov, T. and Repin, D. (2020). Automated deep abstractions for stochastic chemical reaction networks. *ArXiv*, abs/2002.01889.
- Qin, Y., Luo, H., Zhao, F., Fang, Y., Tao, X., and Wang, C. (2023). Spatio-temporal hierarchical mlp network for traffic forecasting. *Information Sciences*, 632:543–554.
- Ramakrishnan, N. and Soni, T. (2018). Network traffic prediction using recurrent neural networks. *2018 17th IEEE International Conference on Machine Learning and Applications (ICMLA)*, pages 187–193.
- Salinas, D., Bohlke-Schneider, M., Callot, L., Medico, R., and Gasthaus, J. (2019). High-dimensional multivariate forecasting with low-rank gaussian copula processes. *ArXiv*, abs/1910.03002.
- Shaheen, S., Chan, N., Bansal, A., and Cohen, A. (2015). Shared mobility: Definitions, industry developments, and early understanding. Technical report, Institute of Transportation Studies, UC Berkeley.
- Shaygan, M., Meese, C., Li, W., Zhao, X. G., and Nejad, M. (2022). Traffic prediction using artificial intelligence: Review of recent advances and emerging opportunities. *Transportation Research Part C: Emerging Technologies*, 145:103921.

Bibliography

- Song, X., Liu, Y., Xue, L., Wang, J., Zhang, J., Wang, J., Jiang, L., and Cheng, Z. (2019). Time-series well performance prediction based on long short-term memory (lstm) neural network model. *Journal of Petroleum Science and Engineering*, 186:106682.
- Strobel, L. and Pruckner, M. (2022). Feasibility of completely electrified two-way car sharing. In *2022 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm)*, pages 206–211.
- TLC Factbook (2024). Tlc factbook. <https://www.nyc.gov/site/tlc/about/data-and-research.page>. Accessed: 2024-07-09.
- Tokuyama, Y., Fukushima, Y., and Yokohira, T. (2018). The effect of using attribute information in network traffic prediction with deep learning. *2018 International Conference on Information and Communication Technology Convergence (ICTC)*, pages 521–525.
- United Nations Environment Programme (2024). Electric mobility. <https://www.unep.org/topics/transport/electric-mobility>. Accessed: 2024-03-05.
- Vazifeh, M., Zhang, H., Santi, P., and Ratti, C. (2015). Optimizing the deployment of electric vehicle charging stations using pervasive mobility data. *ArXiv*, abs/1511.00615.
- Vinayakumar, R., Soman, K., and Poornachandran, P. (2017). Applying deep learning approaches for network traffic prediction. *2017 International Conference on Advances in Computing, Communications and Informatics (ICACCI)*, pages 2353–2358.
- Wang, Liu, Y., fei Yang, Y., Cai, W., xuan Wang, D., and Jia, Z. (2019a). The location of electric vehicle charging stations based on frlm with robust optimization. *Int. J. Pattern Recognit. Artif. Intell.*, 33:1959027:1–1959027:19.
- Wang, N., Tang, L., Zhang, W., and Guo, J. (2019b). How to face the challenges caused by the abolishment of subsidies for electric vehicles in china? *Energy*.
- Wang, S., Bi, S., and Zhang, Y. (2021). Reinforcement learning for real-time pricing and scheduling control in ev charging stations. *IEEE Transactions on Industrial Informatics*, 17:849–859.
- Wu, F. and Sioshansi, R. (2017). 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.
- Xie, R., Wei, W., Wu, Q., Ding, T., and Mei, S. (2020). Optimal service pricing and charging scheduling of an electric vehicle sharing system. *IEEE Transactions on Vehicular Technology*, 69:78–89.
- Ye, J., Sun, L., Du, B., Fu, Y., and Xiong, H. (2021). Coupled layer-wise graph convolution for transportation demand prediction. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35:4617–4625.
- Yin, X., Wu, G., Wei, J., Shen, Y., Qi, H., and Yin, B. (2021). Deep learning on traffic prediction: Methods, analysis, and future directions. *IEEE Transactions on Intelligent Transportation Systems*, 23:4927–4943.
- Yu, G. and Zhang, C. (2004). Switching arima model based forecasting for traffic flow. *2004 IEEE International Conference on Acoustics, Speech, and Signal Processing*, 2:ii–429.

Bibliography

- Zhang, K., Guo, H., Yao, G., Li, C., Zhang, Y., and Wang, W. (2018). Modeling acceptance of electric vehicle sharing based on theory of planned behavior. *Sustainability*.
- Zhao, L., Song, Y., Zhang, C., Liu, Y., Wang, P., Lin, T., Deng, M., and Li, H. (2018). T-gcn: A temporal graph convolutional network for traffic prediction. *IEEE Transactions on Intelligent Transportation Systems*, 21:3848–3858.
- Zhou, X. and Mahmassani, H. (2007). A structural state space model for real-time traffic origin–destination demand estimation and prediction in a day-to-day learning framework. *Transportation Research Part B: Methodological*, 41:823–840.
- Zhu, H., Xie, Y., He, W., Sun, C., Zhu, K., Zhou, G., and Ma, N. (2020). A novel traffic flow forecasting method based on rnn-gcn and brb. *Journal of Advanced Transportation*.
- Çalık, H. and Fortz, B. (2019). A benders decomposition method for locating stations in a one-way electric car sharing system under demand uncertainty. *Transportation Research Part B: Methodological*.

