

Data-Driven Bike-Share Ridership Prediction and Network Optimization

Ghazaleh Mohseni Hosseinabadi

A Thesis Submitted to the Faculty of Graduate Studies

In Partial Fulfillment of the Requirements

for the Degree of Doctor of Philosophy

Graduate Program in Civil Engineering

York University

Toronto, Ontario

©Ghazaleh Mohseni Hosseinabadi, 2025

Abstract

Shared micro-mobility systems, particularly station-based bike-sharing networks, have become key components of urban transportation ecosystems, offering flexible and sustainable solutions for first- and last-mile travel. As these systems expand, their planning and operation require advanced analytical approaches that account for spatial, temporal, and technological complexities. This dissertation addresses critical gaps in existing methodologies by developing integrated models for ridership prediction, station placement optimization, and infrastructure planning in electrified bike-sharing networks.

First, a station-to-station ridership prediction framework is developed using a customized Graph Neural Network (GNN). Leveraging a tailored GraphSAGE architecture, the model integrates network topology, sociodemographic features, and station-level attributes, minimizing weighted prediction errors with a focus on peak-period ridership. Applied to the Toronto system, the GNN outperforms linear, spatial, and tree-based benchmark models, demonstrating its capacity to capture latent network dependencies and support demand-responsive planning.

In the second chapter, a continuum approximation (CA) model is proposed to optimize the spatial configuration of bike-sharing stations. A force-based algorithm balances attraction forces from demand centers and repulsion/attraction forces between stations to determine optimal placements in continuous urban space. This ridership-driven approach departs from conventional accessibility-based methods, directly aligning station locations with demand intensities. The model is validated through numerical experiments and applied to Vancouver, revealing optimal spacing patterns and data-driven strategies for network expansion under varying demand conditions.

To address the operational complexities introduced by electric bikes (e-bikes), the next chapter of this research incorporates battery and charging constraints into infrastructure planning. A Markovian framework models State of Charge (SOC) transitions using a two-dimensional birth-death process that distinguishes high- and low-SOC bikes. A heuristic optimization algorithm is developed to identify charger deployment strategies that maximize expected ridership. A single-pool approximation is introduced to enhance computational scalability. The algorithm

achieves near-optimal solutions and sensitivity analyses reveal how number of chargers, dock capacity, and charging rates affect system performance. Real-world applications in Pittsburgh, Vancouver, and San Francisco illustrate the framework’s ability to identify high-impact charger locations and adapt to diverse electrification scenarios.

Finally, the proposed models are integrated within a web-based, GIS-enabled decision-support tool for interactive scenario analysis to bridge the gap between theoretical development and practical implementation. The tool integrates predictive, prescriptive, and descriptive analytics to enable planners to explore “what-if” scenarios, visualize trade-offs across equity, connectivity, and utilization goals, and support adaptive, data-informed decision-making. The developed tool has been customized for multiple cities, including Toronto and Vancouver, demonstrating its applicability and scalability across diverse urban contexts.

This research advances methodological foundations and practical tools for developing resilient, data-driven, and electrified bike-sharing networks across diverse urban environments.

Keywords: Bike-Sharing Systems, Micro-Mobility, Ridership Prediction, Graph Neural Networks, Continuum Approximation, Facility Location Optimization, Electric Bikes, Charging Infrastructure, State of Charge Dynamics, Markov Models, Interactive Decision-Support Tools.

Acknowledgements

Pursuing this PhD has been far more than an academic journey; it has been a journey of growth, discovery, and learning on multiple levels. I entered this program as a girl with big dreams, and this experience has shaped me in ways I never could have imagined. I have mastered analytical and technical skills, learned to think critically, navigated challenges, and truly embraced the life of a researcher. This path has taught me how to conduct research passionately, waking up each day excited to ask questions, explore ideas, and create something meaningful.

This transformation was only possible because of the extraordinary guidance I received. To my supervisor, Dr. Mehdi Nourinejad, thank you for believing in me, giving me the space to think bigger and beyond boundaries, take risks, and grow. Working with you has been an invaluable opportunity for which I will always be grateful. To Dr. Peter Park, your support, thoughtful advice, and mentorship have been a source of strength throughout this journey. I appreciate the wisdom, patience, and generosity that both of you have shared so openly and feel fortunate to have learned from your examples. I am also grateful to the other members of my committee for their valuable feedback, and to the collaborators in industry whose insights enriched this research.

I owe my deepest thanks to my husband, whose support and willingness to follow me across the world have made all of this possible. You stood by me through every challenge, celebrated every small victory, and reminded me of my strength on the hardest days. Your love, patience, and belief in me have been my anchor and my greatest gift throughout this journey. I am also endlessly grateful to my parents and family, who have been my foundation at every step of my life. Your unwavering support and encouragement, even from afar, have meant everything to me. I am so fortunate to have grown up with your strength behind me.

Finally, I want to thank my friends—those back home and those I've come to know here—as well as the great colleagues who have become true friends along the way. Your kindness, laughter, and support have enriched this journey beyond words.

To all of you: thank you for standing by me throughout this journey. I am profoundly grateful.

Publications and Awards

Chapters of this dissertation have been published, submitted for peer review, or recognized through scholarly awards and professional honors.

- Chapter 2 is based on a manuscript currently under revision at Transportation Research Part A.
- Chapter 3 is based on a manuscript currently under revision at Transportation Research Part C.
- Chapter 4 is based on a manuscript currently under review at Management Science.
- The decision-support tool developed in Chapter 5 was selected as a semi-finalist for the INFORMS Innovative Applications in Analytics Award (IAAA) 2025, recognizing its real-world applicability and impact in transportation analytics. It also received the Innovation Award of the Year (2023) from the Institute of Transportation Engineers (ITE) and was named a finalist in the 2024 Transit Data Challenge organized by the Transportation Research Board (TRB). The tool has been piloted by the City of Vancouver, with additional pilot projects planned in collaboration with other municipalities.

Contents

	Page
Abstract	ii
Acknowledgements	iv
Publications and Awards	v
Contents	vi
List of Tables	x
List of Figures	xi
1 Introduction	1
1.1 Shared Micro-Mobility Systems: Bike-Sharing	1
1.2 Problem Statement	3
1.3 Research Aim and Specific Objectives	4
1.4 Research Scope and Contributions	5
1.5 Thesis Outline	6
2 Ridership Prediction	9
2.1 Introduction	10
2.2 Literature Review	12
2.2.1 Station-Level Prediction	13
2.2.2 Link-Level Prediction	16
2.3 Methodology	18
2.3.1 Problem Statement	18

2.3.2	Model Architecture	18
2.4	Data	24
2.4.1	Data set and Pre-processing	24
2.4.2	Station Features	27
2.5	Analysis	30
2.5.1	Benchmarks	31
2.5.2	Link Weight Prediction for G_{2022}	39
2.5.3	Expansion Design	40
2.6	Conclusions	43
3	Station Placement Optimization	45
3.1	Introduction	46
3.2	Literature Review	49
3.2.1	Bike-Sharing Station Location Optimization	49
3.2.2	Continuum Approximation in Logistics	51
3.3	Problem Statement	54
3.4	Solution Methodology	58
3.4.1	Influence Areas and Constraints	58
3.4.2	A Force-based Algorithm	60
3.4.3	Algorithm Steps	66
3.5	Computational Example	70
3.5.1	Bi-Focal Origin-Destination Demand Analysis	71
3.5.2	Gradient Demand Impact Analysis	80
3.6	Case Study	84
3.7	Conclusion	91
4	Electrification and Charging Infrastructure	94
4.1	Introduction	95
4.2	Literature Review	100
4.2.1	Inventory Management in Regular Bike-Sharing Systems	100

4.2.2	Charging Infrastructure and Battery Management	102
4.3	Methodology	105
4.3.1	Markov Chain Model	106
4.3.2	Optimization Model for Charger Placement	110
4.4	Solution Methodology	115
4.4.1	Ridership Derivation: Equilibrium Computation for a Fixed Charger Configuration	115
4.4.2	Simplification Methods for Ridership Derivation in Large-Scale Problems	117
4.4.3	Ridership Optimization: A Heuristic Search for Charger Configuration	121
4.5	Analysis	122
4.5.1	Station-Level analysis of Markov Chain	122
4.5.2	Computation Example with Exact Probability Model	125
4.5.3	Impact of Number of Chargers	129
4.5.4	Impact of Station Dock Capacity	131
4.5.5	Impact of Charging Rate	133
4.5.6	Comparative Impact of Key System Parameters	135
4.5.7	Comparison of Exact Probability vs. Single-Pooling State Approxima- tion Model	136
4.6	Real Case Studies: Pittsburgh, Vancouver, San Francisco	139
4.7	Conclusion	142
5	Decision-Support Tool Development and Practical Implementation	144
5.1	Overview and Motivation	144
5.2	System Architecture and Data Integration	145
5.3	Core Functionalities and Design	147
5.4	Scalability and Contributions	152
5.5	Case Study: Toronto	153
5.5.1	Scenario Design and Objectives:	154
5.5.2	Optimization and Analytical Workflow	155
5.5.3	Results and Comparative Insights	157

5.5.4	Concluding Reflections, Value, and Future Directions	158
6	Conclusions and Future Work	160
6.1	Summary of the Research	160
6.2	Key Contributions	161
6.3	Policy and Planning Implications	163
6.4	Limitations	164
6.5	Directions for Future Research	165
6.6	Final Remarks	167
	Bibliography	168
	Appendices	180
A	Appendix: List of Notations for Chapter 3	180
B	Appendix: Global balance equations and steady-state solutions for Small C_s . .	182
C	Appendix: Shared E-Scooter Systems	185
C.1	Data and Modeling Framework	185
C.2	Analytical Capabilities and Decision-Support Tool	187
C.3	Policy Implications and Future Directions	188

List of Tables

2.1	POIs Categories	28
2.2	Bike-sharing Network Properties	29
2.3	Hyperparameter Tuning	31
2.4	Models performance comparison for G_{2021} and G_{2022}	33
4.1	Comparison of Optimal Configurations, Ridership, and Computation Time. . .	128
4.2	Comparison of exact probability vs. single-pooling state approximation model .	137
5.1	Summary of Planning Scenarios	155
5.2	Comparative Performance Metrics Across Planning Scenarios	157

List of Figures

1.1	Thesis structure and logical flow of chapters	7
2.1	Architecture of the GNN model	21
2.2	Bike-sharing growth map in Toronto	25
2.3	Toronto 2016-2022, a)ridership, b)monthly trips, c)distance cumulative distribution function (CDF)	26
2.4	Stations' Measure Variations in the Network, a)Nodal strength, b)Weighted clustering coefficient, c)Betweenness centrality	30
2.5	Architecture of the ANN model	33
2.6	Predicted Versus Actual Flows in Neural Network Models: a) ANN (Artificial Neural Network), b) GNN (Graph Neural Network).	34
2.7	Incidence of Negative Predictions in LR, SLR, and XGBoost Models	35
2.8	The relationship between true and predicted flow	36
2.9	a)Comparative Performance Trends of GNN-WMSE and XGBoost (best-performed benchmark model) in Predicting High Flow Values; b)Differential Performance Analysis of GNN-WMSE and XGBoost (best-performed benchmark model) in Predicting High Flow Values	38
2.10	Top-500 accurate link weight predictions versus actual weights- a)2022, b)2021.	40
2.11	Expansion scenarios results	43
3.1	a)Ridership-distance(km) distribution, b)Station-station force concept, c)Station-to-station force function, d)Demand-to-station force function.	63

3.2	Bike-sharing stations, disks, disk's radius, and service areas: a)Infeasible overlapping results in accessibility force, b)Station-to-station forces for station 1 in various scenarios, c)Demand-to-station forces for station 1.	64
3.3	Configuration for an example of $N=20$, a)Origin and destination density heat-maps, b)Initial configuration, c)Station-to-station forces at iteration 1, d)Resultant forces at iteration 1, e)Station movements and updated influence areas at iteration 1.	72
3.4	Station movement and influence area adjustments over nine iterations for an example of $N=20$. Each sub-figure (a) to (i) represents an iteration where stations are re-positioned and their influence areas are recalculated.	73
3.5	a)Total ridership trend, b)Optimized station locations, c)Accessibility heat map	74
3.6	Station movement and influence area adjustments over nine iterations for an example of $N=10$, smaller-scale scenario. a) to e) Represents an iteration where stations are re-positioned and their influence areas are recalculated, f)Station movements tracking over iterations.	75
3.7	Comparative analysis of station distribution and ridership dynamics across θ or D_{opt} values, a) to c) Spatial distribution of bike-sharing stations for varying D_{opt} values, d)A composite view of station distributions for all D_{opt} values, e)Trends in total ridership in relation to different D_{opt} settings.	76
3.8	Comparative analysis of station distribution and ridership dynamics across η values, a) to c) Spatial distribution of bike-sharing stations for varying η Values, d)A composite view of station distributions for all η values, e)Trends in total ridership in relation to different η settings.	78
3.9	Scalability Analysis Results; a)Relationship Between Number of bike-sharing Stations and Iterations to Equilibrium, b)Station locations distribution, $N=50$. .	80
3.10	Gradient demand scenario results, $D_{opt} = 200$, $\eta = 0.6$, a)Origin and destination demand heat maps, b)Initial configuration, b)Station movements and updated influence area in iteration 1, d)Final station locations, c) Total ridership trend.	81

3.11	Comparative analysis of station distribution and ridership dynamics across η values in gradient demand scenario, a) to c) Spatial distribution of bike-sharing stations for varying η Values, d)A composite view of station distributions for all η values, e)Trends in total ridership in relation to different η settings.	84
3.12	a)City of Vancouver bike-sharing system, b)Origin demand density, c)Destination demand density heat map, d)Combined origin and destination demand heat map.	86
3.13	Five new station scenario, a)Resultant forces in iteration 1, b)Station movements in iteration 1, c)Final location of new stations.	88
3.14	a)Station movements-10 new station- Iteration 1, b)Station movements-20 new station- Iteration 1, c)Optimized location of 10 new stations, d)Optimized location of 20 new stations.	89
3.15	Optimized station configuration for the full deployment scenario, where all stations are placed simultaneously without any fixed infrastructure.	90
4.1	State transition diagram for a station with capacity $C_s = 3$. State (i, j) represents i low- and j high-SOC bikes.	108
4.2	Birth-death process representations for simplified 1D Markov chain, a)Single-pooling state approximation, b)Instant charging approximation	118
4.3	a) State transition diagram($C_s = 2$), station-Level analysis of steady state probabilities versus: b) Charging rate, c) Low-SOC rental rates, d) High-SOC arrival rate	123
4.4	Station-level Markov Chain analysis under different policies: (a–c) State transition diagrams, d) Steady-state probability vs. charging rate for first policy ($\lambda^H = 0$), e) Steady-state probability vs. charging rate for first policy ($\mu^L = 0$), f) Steady-state probability vs. charging rate for first policy ($\lambda^H = \mu^L = 0$).	126
4.5	a) Station locations and O-D distribution, where edge thickness represents trip frequencies, b) Convergence of arrival rates across stations over iterations, confirming the attainment of equilibrium, c) Ridership trend across iterations for different approaches.	127

4.6	Number of chargers sensitivity analysis results. a) Total ridership and optimal ridership, b) Marginal ridership gain, c) Computation runtime, d) Fraction of high-SOC bikes for each station, e) Probability of being at capacity for each station, f) Optimal charger configurations.	130
4.7	Dock capacity sensitivity analysis results. a) Total ridership and optimal ridership, b) Marginal ridership gain, c) Computation runtime, d) Fraction of high-SOC bikes for each station, e) Probability of being at full capacity for each station, f) Optimal charger configurations.	132
4.8	Charging rate sensitivity analysis results. a) Total ridership, b) Marginal ridership gain, c) Computation runtime, d) Fraction of high-SOC bikes for each station, e) Probability of being at full capacity for each station, f) Optimal charger configurations.	134
4.9	Comparison of changes in key system parameters (number of chargers, dock capacity, charging rate) on: (a) Total ridership, (b) Marginal ridership gain, and (c) Computation time.	136
4.10	Models comparison in terms of station capacity and a)Computation time, b) Ridership	138
4.11	Comparison of steady-state probabilities in single-pooling approximation model and exact probability model for a network of 10 stations.	139
4.12	Pittsburgh bike-sharing network: Optimal charging-station distributions for approximately 10, 30, 50 percent electrification.	140
4.13	Vancouver and San Francisco bike-sharing networks: (a) and (b) optimal charging-station configurations in Vancouver for approximately 10% and 50% station electrification; (c) and (d) optimal charging-station configurations in San Francisco for approximately 10% and 50% station electrification.	141
5.1	System architecture of BikesharePro: from data sources to deployment.	147

5.2	Snapshot of the tool calibrated for Toronto. Hexagon fills represent spatial coverage and station presence, while circle overlays depict station-level outflow (green) and inflow (red), scaled by volume. The underlying blue hex grid encodes demographic and accessibility indicators.	148
5.3	Descriptive analytics capabilities example of the tool demonstrated through selected panels. a) OD ridership shown as volume-colored lines, filterable by year and season, b) Transit integration heatmap: red areas indicate no bike-share station within a walkable distance of a transit stop, yellow areas indicate one station, and blue areas indicate two or more, c) Cycling congestion on bike-friendly routes based on ridership and routing API, d) Station-level view with points of interest and monthly bike availability, e) Net flow visualization at station level with shaded circles—green for net outflow, red for net inflow—and corresponding system-level ridership metrics in the lower panel, f) Cost and revenue inputs with financial summaries for system modifications.	149
5.4	Interactive data editing panels. a) Cell-level demographics and land use editor, b) Station property editor with capacity, charger availability, and status settings.	150
5.5	Predictive analytics for station deployment. a) Adding a new station with configurable attributes such as capacity, installation year, and charging capability, b) Predicted OD flows from the newly added station visualized as directional lines.	150
5.6	Prescriptive station placement. Both figures use the same weights for accessibility, transit integration, and equity. Red hexagons indicate newly recommended station locations. a) With low connectivity importance, the model favors outer, underserved regions, b) With higher connectivity importance, the model prioritizes integration with the existing network, placing stations closer to current infrastructure.	151

5.7	Deployment of BikesharePro across diverse urban environments. It has been applied using real operational data in Vancouver, Boston, San Francisco, Chicago, and Cincinnati, and evaluated with proxy data and similar models in Hamilton, Ottawa, and Mississauga.	152
5.8	Scenario-specific spatial weighting distributions and corresponding optimized station locations. Figures (a), (c), and (e) display the weighted planning maps for Scenario 1 (Ridership-focused), Scenario 2 (Transit Integration), and Scenario 3 (Equity-Prioritized), respectively. Darker blue shading indicates higher priority based on the weightings. Figures (b), (d), and (f) show the resulting optimized new station locations, highlighted in red, as generated by the Prescriptive Optimization Module.	156
1	Main map interface showing pickup volumes by hexagonal cell overlaid with population background shading. The sidebar allows users to filter by year, season, demographic layer, POI, and operational zones.	187
2	Prediction configuration panels. (a) and (b) Daily and seasonal prediction inputs. (c) Time-lapse viewer for exploring temporal demand dynamics.	188
3	Various system-wide analytics interfaces. Visualizations include seasonal pickup/drop-off totals, demand distributions by zone, and transit stop-level performance indicators.	188
4	Transit integration analysis based on: (a) transit stops, (b) transit hubs with high ridership volume.	189

Chapter 1

Introduction

1.1 Shared Micro-Mobility Systems: Bike-Sharing

Shared micro-mobility systems are urban transportation services that provide short-term access to lightweight, low-speed vehicles, such as bicycles, electric bicycles (e-bikes), and scooters. These systems have emerged as a flexible and sustainable mobility option, offering an alternative to private car use and enhancing connectivity across multimodal transportation networks. Users typically access shared vehicles through mobile applications or kiosks, enabling point-to-point travel without the need for personal ownership and aligning with urban goals for accessibility, affordability, and environmental sustainability.

Station-based bike-sharing has become one of the most structured and widely implemented systems across cities among the various shared micro-mobility modes. In this model, users rent and return bicycles at fixed docking stations distributed throughout the service area, allowing for controlled vehicle distribution and reliable access points. As of 2023, more than 3,000 cities in 92 countries operate bike-sharing programs (Meddin et al., 2022), reflecting their widespread adoption and institutionalization within urban transportation systems. Well-known systems such as Citi Bike in New York, BIXI in Montreal, and Mobi in Vancouver exemplify the scalability and versatility of station-based networks. These systems offer independent mobility options and play a critical role in enhancing first- and last-mile connectivity by linking residential neighborhoods to transit hubs, employment centers, and other key destinations.

A defining feature of station-based bike-sharing systems is their reliance on strategically placed docking stations, which form the structural foundation of the network. These stations serve as user access points and spatial anchors that shape travel behavior and define the system's service coverage. The location and configuration of stations influence origin-destination (OD)

ridership flows, determining how users move across the network and where demand is concentrated. Accurately predicting these flows is critical for tactical-level planning, as it informs decisions about where to place new stations, how to allocate capacity, and how to align network growth with evolving mobility needs. While operational concerns such as vehicle rebalancing and availability management are also important, they are fundamentally shaped by upstream planning decisions. As such, station design is key in aligning land use, demographic patterns, and multimodal connectivity with system performance, ridership outcomes, and accessibility goals.

In recent years, integrating e-bikes into bike-sharing systems has reshaped the landscape of shared urban mobility. E-bikes offer battery-assisted propulsion to reduce physical barriers and extend feasible travel distances. E-bikes have broadened the appeal of shared cycling by attracting a more diverse user base, enabling longer trips, and enhancing accessibility in topographically or physically challenging environments. Their increasing popularity (National Association of City Transportation Officials, 2024) has significantly influenced ridership levels and system utilization patterns, reinforcing their role as both a complement to and a substitute for conventional bicycles. E-bike adoption introduces new operational considerations related to battery life and recharging requirements. Multiple charging strategies have been introduced to support this transition, including plug-in charging stations, battery-swapping systems, and hybrid approaches that integrate multiple power sources. In station-based systems, the spatial deployment of charging infrastructure—alongside the growing share of electric vehicles in the fleet—has introduced new system design layers that intersect with planning and operations. As station-based systems continue to expand and evolve, electrification has become a defining feature of modern bike-sharing networks, with implications for network layout, accessibility, and long-term sustainability.

The continued evolution of shared micro-mobility, particularly in its station-based and electrified forms, has reshaped how cities plan transportation. Bike-sharing systems are now integral to advancing sustainability, equity, and multimodal connectivity goals. This shift has prompted growing academic and institutional attention toward optimizing these systems' design, expansion, and integration within the broader urban mobility framework. As their influence expands,

so does the need for data-driven, scalable, and interdisciplinary approaches to guide their development in diverse urban contexts.

1.2 Problem Statement

Bike-share systems are typically introduced on a small scale as a pilot and gradually expanded in response to demand growth. For example, New York City launched its CitiBike program 2013 with 329 stations. By the end of 2022, it had expanded its coverage area and number of stations to 750 in response to the growing interest and usage (Cit (2023)). The ongoing expansion of the bike-share network can be more impactful by optimizing the system’s design and efficiency over time. The success of bike-share systems relies on data-driven optimization, ensuring effective planning for long-term sustainability. Planning a bike-share network involves station deployment, deactivation, re-deployment, dock sizing, and fleet sizing.

Demand Prediction Challenges: An important aspect of bike-share systems expansion planning is predicting potential ridership between new and existing stations by analyzing origin and destination patterns. Such predictions assist in making informed decisions about station placement, dock allocation, and capacity adjustments and help to pinpoint areas that need targeted interventions to manage demand and optimize network infrastructure. While inflow/outflow models focusing on station-level prediction have been extensively studied, relatively little attention has been paid to station-to-station ridership prediction, especially network-sensitive methods that capture inter-station dependencies. As a result, many existing approaches fail to account for the structural properties of the network, which play a significant role in shaping usage patterns.

Station Placement Limitations: The spatial configuration of the stations further complicates the planning process. Optimal station placement must balance local accessibility with broader network connectivity to maximize ridership. Although discrete facility location models—typically based on mixed-integer programming—have been widely applied, they often require extensive location-specific data and are computationally intensive for large urban networks. Moreover, these models rarely incorporate network-wide ridership dynamics, limiting their ability to maximize system usage at scale.

Electrification Complexity: The increasing integration of electric bikes (e-bikes) further complicates the planning process. E-bikes offer substantial benefits in terms of trip length, accessibility, and user adoption, but they introduce unique operational constraints related to battery depletion and recharging. This creates new requirements for infrastructure planning, including the spatial allocation of charging stations and the modeling of SOC dynamics across the network. Existing research has examined charging logistics and rebalancing strategies, but few studies incorporate SOC-aware modeling into system-level optimization frameworks that also consider ridership impacts.

There is thus a critical need for an integrated planning framework that:

- Predicts OD ridership using network-sensitive methods;
- Optimizes station placement in continuous space based on ridership potential and connectivity; and
- Incorporates SOC constraints and charging infrastructure into electrified network planning.

This research responds to these needs by developing analytical and decision-support models that enable scalable, data-driven, and electrification-aware planning for bike-sharing systems.

1.3 Research Aim and Specific Objectives

The main goal of this study is to develop a data-driven, integrated planning framework that advances the modeling and optimization of station placement, ridership prediction, and electrification strategies within shared cycling networks. This framework also supports more informed and adaptive policy decisions related to station planning, fleet management, infrastructure expansion, and multimodal integration. The specific objectives are outlined as follows:

1. **Predict station-to-station ridership flows:** Design a machine learning model to estimate origin-destination demand using network topology, station attributes, and contextual variables. The model captures spatial dependencies and improves prediction accuracy compared to traditional station-level approaches.

2. **Optimize station placement in continuous space:** Formulate an optimization algorithm that repositions stations based on attraction-repulsion forces to maximize ridership, while accounting for inter-station dynamics and spatial demand coverage.
3. **Integrate e-bike charging constraints into system planning:** Develop a scalable model incorporating SOC dynamics to support charger placement decisions and enhance ridership and operational efficiency in electrified bike-sharing systems.
4. **Validate the framework across diverse real-world systems:** Apply the integrated models to case studies of varying scale and complexity. Assess model performance, adaptability, and practical relevance to demonstrate the framework’s generalizability and planning value.
5. **Implement a GIS-based decision-support tool:** Develop an interactive web platform that operationalizes the research findings. The tool enables scenario-based planning and trade-off analysis to support urban planners and policymakers’ strategic, equity-oriented, and transit-integrated decision-making.

1.4 Research Scope and Contributions

This research focuses on strategic planning for station-based bike-sharing systems, specifically addressing medium- to long-term decisions related to network design, infrastructure expansion, and electrification. It does not cover operational aspects such as vehicle rebalancing, real-time fleet management, or short-term demand forecasting.

The core contribution of this work is the development of an integrated analytical and decision-support framework tailored to the unique spatial and infrastructural challenges of shared micro-mobility. Methodologically, the research advances the state of the art in three key areas.

First, it introduces a customized Graph Neural Network (GNN) model to predict origin-destination flows between bike-sharing stations. The model leverages GraphSAGE architecture to embed station-level features within the network topology, capturing spatial and structural dependencies. By focusing on peak-period ridership, the model improves planning relevance for high-demand periods and outperforms traditional inflow/outflow-based approaches.

Second, a continuous approximation (CA) framework featuring a force-based algorithm is proposed for optimizing bike-sharing station locations in continuous space, to maximize ridership. Unlike traditional accessibility-based models or discrete optimization methods, which can become computationally challenging at large scales, this approach effectively captures network-level ridership effects. It models attraction forces from spatially distributed demand hotspots and repulsion/attraction forces to maintain suitable spacing between stations. The algorithm allows stations to move freely within the service area, dynamically adjusting to changes in demand densities and origin-destination patterns.

Third, the thesis advances electrification-aware planning by proposing a Markovian modeling framework that integrates SOC dynamics, charger placement, and ridership optimization for station-based e-bike networks. The system is modeled as a two-dimensional birth-and-death process that tracks low- and high-SOC bikes under charging transitions. Steady-state probabilities are derived to estimate expected ridership, forming the objective function of a station-level charger allocation problem. A scalable heuristic with a single-pool approximation is developed to solve this problem efficiently, achieving near-optimal results while reducing computational time.

This thesis also delivers a web-based, GIS-enabled planning tool that integrates predictive, prescriptive, and interactive components to bridge research and practice. Designed for deployment in real-world planning environments, the tool has been customized for and tested in multiple cities, such as Toronto, Vancouver, and San Francisco, demonstrating its scalability and practical utility. This thesis contributes methodological innovations and a deployable, city-tested decision-support system for the scalable and sustainable planning of bike-sharing networks.

1.5 Thesis Outline

This dissertation is organized into six chapters, reflecting the progressive development of an integrated planning framework for bike-sharing systems. The structure moves from problem formulation and demand prediction to network optimization, electrification modeling, and real-world deployment. The logical relationships between chapters and their respective contributions to the overall framework are illustrated in Figure 1.1. Chapters 2 to 4 are not directly fed into one

another; instead, they are modular but conceptually aligned, with their integration demonstrated in Chapter 5.




Figure 1.1: Thesis structure and logical flow of chapters

- **Chapter 1 – Introduction:** Establishes the research context, defines the problem, outlines objectives, contributions, and scope, and presents an overview of the thesis structure.
- **Chapter 2 – Graph-Based Ridership Prediction:** Develops a GNN framework to predict station-to-station ridership flows, capturing spatial dependencies within the bike-sharing network.
- **Chapter 3 – Station Placement Optimization:** Introduces a force-based CA model to optimize station locations in continuous space, aiming to maximize ridership while preserving spatial coverage and connectivity.
- **Chapter 4 – Electrification and Charging Infrastructure:** Proposes a Markovian planning model that incorporates SOC dynamics and charger placement into network-level

optimization for e-bike sharing systems.

- **Chapter 5 – Decision-Support Tool Development and Practical Implementation:** Describes the development and deployment of a GIS-based planning tool, integrating description, prediction, and optimization components for real-world decision-making.
- **Chapter 6 – Conclusions and Future Work:** Synthesizes the main findings, highlights broader contributions, discusses limitations, and proposes future research and development directions.
- **Appendices:** The appendices include supporting technical content and a related extension study. Appendix A provides a list of notations used in Chapter 3. Appendix B outlines the global balance equations and steady-state solutions for small C_s values in the Markovian electrification model presented in Chapter 4. Appendix C presents a complementary study on hybrid shared e-scooter systems, illustrating the adaptability of the dissertation's methodology to operational micromobility planning.

Chapter 2

Ridership Prediction

Abstract: Ridership prediction in station-based bike-sharing services improves station planning, fleet management, and network design. Ridership inflow and outflow prediction at the station level has received significant attention through trip production and attraction models. However, station-to-station ridership has been studied less, despite its widespread applications in use cases such as bike-lane planning or fleet electrification. This study introduces a Graph Neural Network (GNN) to model station-to-station ridership using a customized Graph Sample and Aggregate framework to generate node embeddings and minimize the weighted Mean Squared Error for peak periods. The model incorporates the characteristics of the network, sociodemographic features, and station properties. We present the case study of Bikeshare Toronto to train and test the GNN model and benchmark it against other standard prediction methods. We show that the GNN outperforms linear regression, spatial regression, XGBoost, and artificial neural networks due to its ability to capture the impact of the network structure on ridership patterns. We incorporate the GNN model in five design scenarios focusing on urban core connectivity, suburban access, transit integration, equitable accessibility, and tourist hubs. Each scenario is strategically developed to prioritize and address unique urban challenges. To enhance the model’s application in real-world planning, we embedded the model in a web-based tool for the Cities of Vancouver and Toronto, allowing for further “what-if” scenario analysis in bike-sharing network planning.

Keywords: Bike-sharing, Graph Neural Network, Ridership Prediction, Link Weight Prediction, Micro-mobility.

2.1 Introduction

Bike-sharing is now offered in more than 3,000 cities across 92 countries (Meddin et al., 2022) and has become a prominent mode of active transportation. Many bike-sharing services start as small-scale pilot programs and are progressively expanded to accommodate increasing demand. CitiBike, for example, was launched in 2013 in New York with 329 stations and expanded to 750 stations by 2022 (Cit, 2023). Planning a station-based bike-sharing network involves determining station locations, their dock capacity, and the size and composition of the fleet. The efficiency of network expansion depends on the ability to predict the potential ridership between new and existing stations. Ridership prediction enables planners to make strategic network expansion decisions, such as placing new stations in locations that generate high ridership, while considering the socioeconomic impacts of accessibility to those stations (Liu et al., 2017a).

Existing station-based ridership prediction models use methods of statistical modeling, Machine Learning, and Deep Learning, considering station properties, geographic features, and demographic characteristics (Thomas et al., 2009; Buck and Buehler, 2012; Frade and Ribeiro, 2014; Zhan et al., 2023). Traditional Machine Learning models rely on simplifying assumptions, such as fixed feature sets, which can struggle to address nonlinear relationships in the network structure of bike-sharing trip patterns. Recent developments in Deep Learning and Graph Theory, such as Deep Neural Networks, have introduced new ways of predicting system evolution. Graph Neural Networks (GNN) address these caveats by capturing the network impacts of ridership prediction, where each node is represented by its features and those of its neighboring nodes. This approach captures complex dependencies in large-scale graph data, such as bike-sharing systems, where demand and supply at one station can influence others.

Inflow and outflow ridership prediction at the station level, i.e., trip generation and attraction, has received significant attention since it can provide performance measures such as bike availability at each station (Kim et al., 2019; Guo et al., 2019). Trip production refers to the rate of trips originating from a station, while attraction refers to the rate of trips ending at a station. Such models can improve the balance of fleet distribution in the network to decrease bike shortages at high-demand stations and increase the availability of free docks for returns. Since

trip generation and attraction models do not capture trip distribution, they require secondary models to capture station-to-station ridership. In contrast, station-to-station ridership prediction requires identifying links between station pairs and the corresponding ridership, thus eliminating the need for additional trip distribution models.

Ridership prediction studies that use GNN-based models focus on binary classifications and link strengths as binary representations of ridership (Xiao et al., 2022, 2023). These studies overlook the directed weighted nature of bike-sharing networks and do not offer ridership predictions as continuous variables. Some of such concerns were raised and approached by Müller (2023), who investigated the prediction of flight routes (link presence) and frequencies (link weights) using GNN models, highlighting the potential of graph-based prediction models in such transportation networks. Yang et al. (2020) developed a graph structural model for short-term ridership prediction in bike-sharing systems, emphasizing the importance of incorporating graph structural attributes, such as Out-strength, In-strength, Out-degree, In-degree, and PageRank. The model outperforms traditional benchmarks relying on commonly used features like meteorological data by capturing time-lagged graph variables that describe flow interactions and network dynamics. Despite these advancements, a critical gap remains in treating bike-sharing networks as directed weighted graphs. Addressing this gap forms the core research question of this chapter: What is the added value of capturing the network effects in ridership prediction for tactical-level planning of bike-sharing services?

This study represents each bike-sharing station as a node and defines “flow” to describe the ridership between these nodes. We focus on predicting flows in a weighted directed bike-sharing network using local structural information from the neighbors associated with each station. The model predicts ridership as continuous link weights, addressing the limitations of previous studies that focused on binary classifications. The model is applied to the Toronto bike-sharing network for two consecutive years to evaluate its performance. Results show that the model outperforms the benchmarks, averaging an 18% improvement in RMSE and a 12% improvement in MAE for existing stations. For new stations in the following year, an average improvement of 15% in RMSE and 20% in MAE was achieved. The findings highlight the value of incorporating network effects into predictive models, supporting more data-driven and

efficient decision-making processes in bike-sharing systems. The main contributions include capturing ridership patterns, emphasizing link weight prediction over binary link existence, and developing an interactive GIS-based tool for evaluating what-if scenarios and policy decisions. As a policy analysis and design exercise, we embedded the prediction model in a GIS-based interactive tool to analyze “what-if” scenarios in the expansion plan. The tool captures network dynamics and supports data-driven decision-making for network expansion. The trigger for executing the GNN model in the tool is the addition or removal of stations, upon which the ridership between selected station pairs is re-calculated. Other design features of the network, such as network capacity, which is a significant factor in the GNN prediction model, can also be adjusted, upon which the ridership is re-calculated. The implemented case studies of the GIS tool are the cities of Toronto and Vancouver.

The remainder of this chapter is organized as follows. Section 2.2 reviews relevant literature on bike-share ridership modeling, highlighting key methodological trends and gaps that motivate our approach. Section 2.3 presents the methodological framework, including feature selection, data preprocessing, and model formulation. Section 2.4 details the application of the predictive model to the Toronto bike-sharing network and analyzes the model’s performance across diverse temporal and spatial contexts. Section 2.5 introduces a GIS-integrated scenario analysis tool developed for evaluating network expansion strategies and presents results from simulated deployment scenarios. Finally, Section 2.6 synthesizes the main insights, discusses limitations, and identifies avenues for future research.

2.2 Literature Review

This section reviews the literature on predicting demand and ridership in bike-sharing systems. We categorize the studies into two main areas: station-level prediction and link-level prediction. This categorization facilitates a thorough examination of the methodologies used to predict bike-sharing dynamics, showcasing both traditional and advanced approaches. By delineating these areas, we aim to identify the strengths and limitations of current models, highlight existing research gaps, and state the contributions of this chapter.

2.2.1 Station-Level Prediction

Early studies on bike-sharing demand prediction include linear regression (Thomas et al., 2009; Buck and Buehler, 2012; Frade and Ribeiro, 2014) and multi-level mixed regression models (Rixey, 2013). These approaches were later advanced to more complex Machine Learning techniques, such as gradient boosting machine (Regue and Recker, 2014), clustering methods combined with Bayesian network analysis (Froehlich et al., 2009), or negative binomial models (Hosseinzadeh et al., 2021), to capture the relationship between the built environment (factors like weather conditions, temporal variables such as time of day and week, or proximity to special events) and bike-sharing flows. Some studies integrated clustering with regression-based modeling to account for station-level heterogeneity; for instance, Hyland et al. (2018) proposed a hybrid cluster-regression approach, where bike-sharing stations are first grouped based on trip characteristics using k-means or fuzzy c-means clustering, followed by a multi-level mixed regression model incorporating station-cluster interactions. Preliminary studies identified a combination of factors such as spatial patterns, station proximity effects, sociodemographic factors, and nearby points of interest (POIs) as key influences on bike-sharing usage, with demand often exhibiting nonlinear relationships.

More advanced Machine Learning methods enhance traditional models by leveraging spatial-temporal dependencies, integrating diverse data types, and addressing the inherent complexities of bike-sharing systems. For example, Zhu et al. (2022) applied the Gaussian Process and implemented a grid input structure to capture intrinsic spatial-temporal correlations. They used this approach to model the number of pickups, returns, and idle bikes in New York City’s Citi Bike service. Yang et al. (2019) used a closed-loop simulation-based prediction framework to predict demand at station and city levels while considering spatiotemporal demand shifts between stations. Guidon et al. (2020) also used a tree-based method to predict demand for electric bike-sharing in new cities and to identify the most important predictors of using electric bikes. To better capture spatial dependencies, Dey et al. (2021) integrated origin-destination (OD) flows into station-level demand models and proposed a hybrid approach combining a linear mixed model for weekly demand estimation with a Multiple Discrete Continuous Extreme Value for trip distribution. Their study, using CitiBike data from New York City, highlights the

influence of station capacity, bike lane density, transit accessibility, and land use mix on bike-sharing demand and OD flows. The findings refine spatial demand modeling and emphasize the need for approaches that explicitly capture station interactions. Other studies cluster stations into functional zones based on points of interest and geographical distances (Liu et al., 2017a) and assess the impact of bike-sharing network spatial structures on expansion plans (Kou and Cai, 2021). Additionally, prediction tasks in bike-sharing systems use methods such as pattern recognition models (Sohrabi et al., 2020), sinusoidal models to capture seasonal bicycle usage patterns (Fournier et al., 2017), and methods that address the stochastic nature of bike-sharing systems (Xu et al., 2023). These methods rely on features like the built environment and demand uncertainty. Despite significant advancements in this area, classical Machine Learning models consider many simplified assumptions, such as fixed feature sets, which require more adaptability to changing conditions. This rigidity can lead to lower predictive performance, especially in bike-sharing systems where demand patterns vary significantly across locations. Recently, several studies have applied deep learning methods to capture complex travel patterns. Deep Neural Network models for shared mobility systems can include Feedforward Neural Networks, Recurrent Neural Networks, Convolutional Neural Networks, and GNNs (Jiang, 2022). Liu et al. (2015) developed a bike-sharing station site allocation model using an Artificial Neural Network, considering factors such as usage, required re-balancing efforts, and the public transportation network. Among the Deep Neural Network variants, GNNs stand out as they can handle graph-structured relations. As a result, various disciplines of transportation, including but not limited to traffic forecasting, bike-sharing demand prediction, passenger flow prediction, and multi-modal demand prediction studies (Rahmani et al., 2023) have increasingly adopted GNNs due to their ability to model the relational data and spatial dependencies inherent in transportation networks. GNNs excel by directly incorporating the network topology into their analysis, enabling more accurate predictions that reflect the complexities of mobility patterns and interactions. The fundamental concept of a GNN is that each node is represented by its features and those of its neighbors, known as node embedding. Upon obtaining node representations, various tasks, such as classification, clustering, and prediction, can be performed at the node, link, or graph level (Zhou et al., 2020).

In graph-structured prediction tasks, the structure of the adjacency matrix—which encodes the relationships between nodes—plays a key role in determining model performance. Lin et al. (2018) addressed this by proposing a data-driven Graph Filter model within a Graph Convolutional Neural Network, aiming to learn the adjacency matrix directly from the data. They provided four types of adjacency matrices, including the spatial distance matrix, demand matrix, average trip duration matrix, and demand correlation matrix. Their model captured pairwise correlations by learning hidden heterogeneous relationships between stations to predict station-level hourly demand. Ma et al. (2022) is another study that presented improvements in prediction performance using a spatiotemporal graph attentional long-term short-term memory model to learn the adjacency matrix while incorporating land-use, weather, and user personal information data. The long-term short-term memory neural network was also used to predict the bike-sharing trip and mobility patterns in a dockless bike-sharing network (Xu et al., 2018). A multi-graph Convolutional Neural Network model was proposed to predict demands at the node level by Chai et al. (2018). They introduced a weighted bike-sharing graph, where the weights represented the strengths of relationships between stations. The study considered a weighted bike-sharing graph with the weights representing the station relation strengths. The study reported a 25% reduction in prediction error when comparing their model to base models, including the long short-term memory model, Auto-Regressive Integrated Moving Average and its seasonal version, Gradient Boosting Regression Tree, and Historical Mean.

Predicting bike-sharing demand for pickups and returns using a spatiotemporal Graph Convolutional Neural Network (Xiao et al., 2021) and using GNN to generate embedding for station-based bike predictions (Li et al., 2022) also represented the relationship between each station's demand and its nearby stations in addition to time slots and the dynamic supply patterns between stations. In a recent study, Liang et al. (2023) predicted the monthly inflow and outflow trip generation for bike-sharing expansion using the GNN framework. Their research focused on predicting station-level demand using urban built environment data, spatial interactions, and temporal features. Spatial correlation weights between stations were acquired using attention mechanisms. They demonstrated superior performance by comparing existing methods and alternative deep learning approaches and highlighted the effectiveness of the GNN-based ap-

proach. Another study by Ding et al. (2022) examined the impact of dockless bike-sharing systems using the intervention response module and GNN to analyze the dynamic effects of demand at individual stations. They focused on historical usage data from the London Cycle Hire scheme, finding that the influence of station-less bike-sharing on demand increases over time, highlighting the interdependence between dockless and docked bike-sharing systems.

The literature reviewed above has primarily focused on deep Machine Learning methods for predicting station-level dynamics in bike-sharing networks. While these models provide insights into localized demand and supply, they often overlook inter-station interactions required for optimizing the entire network, such as origin-destination relationships that inform trip flows and redistribution strategies. In contrast, link-level predictions focus on flows between stations, capturing the bike-sharing network’s directed and weighted nature and offering a system-wide perspective. In the next subsection, we discuss the limitations of existing link-level predictions, which are often restricted to binary link presence or simplistic link weights. We then present advancements in link-level prediction that address these gaps and emphasize their significance for optimizing network design.

2.2.2 Link-Level Prediction

Ridership prediction at the link level can be divided into link presence and link weight prediction. Link presence prediction aims to capture the existence of links between station pairs and is commonly treated as a binary classification problem. In contrast, link weight prediction seeks the strength of relationships between stations, i.e., their ridership. Many link prediction algorithms focus on probabilistic, learning-based, and similarity-based approaches (Kumar et al., 2016; Müller, 2023). Tavares et al. (2022) applied GNN to a link prediction task to predict bike station demand. Their approach involved the proposal of a station clustering algorithm based on geographic distances and trip patterns. Xiao et al. (2022, 2023) investigated the impact of network information on bike-sharing ridership prediction using the concept of embedding through Graph Sample and Aggregated (GraphSAGE), a Machine Learning technique that learns node embeddings. GraphSAGE is designed to generate embeddings by sampling and aggregating features from a node’s neighborhood, making it effective for graph-based data. In their initial

study, they presented a GNN-based binary classification model to predict the existence of links in a directed unweighted network. Subsequently, they enhanced their model by incorporating link strengths into the link prediction framework. They systematically varied the link strength to evaluate the model's performance, corresponding to large and small network sizes. The findings suggested that although the current prediction of new stations may not be satisfactory, achieving a more accurate approximation of the adjacency matrix could significantly improve the model's performance.

Most GNN models have been applied in the link prediction context, while only a few have considered link weight prediction. Müller (2023) uses both link presence and weight prediction to predict the potential formation of flights between airports and predict their frequency. The study presented a supervised learning framework based on RNN and evaluated the performance of the proposed model against commonly used similarity-based algorithms. The study predicted the weight of top-L predicted links and calculated RMSE solely on these links. Their findings highlight the challenges of predicting links and edge weights in sparsely connected networks. They conclude that further optimization of the prediction models is necessary before they can be effectively used to inform concrete policy decisions.

The literature on GNN models for bike-sharing prediction focuses on station-level models, often disregarding the significance of link-level predictions in directed weighted graphs that capture the origin and destination of trips and their respective ridership. To reduce overcrowding in popular areas, it is important to optimize the current system, adding more stations or shutting down certain ones. This optimization can be achieved more appropriately by examining the network as a directed weighted graph in a link-level approach while extending the prediction horizon to include long-term trends. Therefore, the primary objective of this chapter is to develop a GNN model that incorporates network properties and neighborhood information within a weighted directed network framework. By predicting the flow between stations in bike-sharing systems, the proposed model supports decision-making in the design and expansion for more data-driven decision-making.

2.3 Methodology

We present the problem statement and the architecture of the GNN model to predict station-to-station ridership. The GNN serves as the base model for assessing the impacts of adding new stations in subsequent sections.

2.3.1 Problem Statement

We express the problem as a link weight prediction task within the context of a weighted directed network, represented as $G(V, E, W)$, where V is the node set (bike-sharing stations), E is the link set representing the existence of flow between a station pair, and W denotes the link weights representing ridership on a link that exists in E . Specifically, for each directed link $(u, v) \in E$, the weight $w_{uv} > 0$ signifies the flow from station u to station v . Given that the network is directed, the weights w_{uv} and w_{vu} can differ, reflecting asymmetric flows between station pairs. The adjacency matrix $A = \{a_{uv}\}$ captures the connectivity of the network, where:

$$a_{uv} = \begin{cases} 1 & \text{if there is a directed link from node } u \text{ to node } v \\ 0 & \text{Otherwise,} \end{cases} \quad (2.1)$$

One possible way of establishing the link connections is a distance (or travel-time) threshold approach whereby two nodes are connected if their distance is less than some threshold δ .

The network structure is the foundational graph of the proposed GNN model. This structure enables the extraction of network properties as node features and facilitates the embedding of station features for prediction. In the subsequent section, we provide a detailed explanation of the architecture of the GNN model.

2.3.2 Model Architecture

The principle of GNNs involves characterizing each node based on its features and the features of its neighboring nodes. This representation, known as node embedding, leverages the structural information of the graph and learns informative node representations that can be utilized for various downstream tasks.

This study uses the GraphSAGE framework to generate node embedding within the GNN model. This framework involves two essential steps. Firstly, a fixed number of neighboring nodes are randomly selected for each node in the network, providing a localized view of the graph. Subsequently, the features of the sampled neighbors are combined to compute an updated representation for the target node. GraphSAGE is an inductive framework with a representation of unseen nodes and adaptation to network changes, such as the addition or removal of nodes. A more detailed description of the algorithm can be found in Hamilton et al. (2017)). GraphSAGE is a suitable choice as it can predict ridership for new stations by leveraging learned neighborhood aggregation functions to infer embeddings for nodes not presented during training. Moreover, it is robust in handling graph evolution and scalable in processing large networks.

An alternative approach for capturing relationships within graph structures is Graph Attention Networks, which uses an attention mechanism to assign different weights to neighboring nodes based on their relative importance. This attention mechanism can be particularly useful for adapting to variable dependencies between nodes. However, Graph Attention Networks operate in a transductive setting, requiring the entire graph structure to be present during training, which limits its flexibility in handling dynamic networks. Given the evolving nature of bike-sharing systems, where new stations may frequently be added, we require an inductive model that can generalize to unseen nodes without retraining from scratch. To achieve a similar effect to attention while maintaining scalability and adaptability, we incorporate a Gaussian-based weighting function within the GraphSAGE framework. The Gaussian distribution is characterized by a bell-shaped curve, which intuitively models the varying influence between stations based on distance. The peak of the Gaussian function represents an optimal range of distances where bike-sharing trips are most frequent - distances that are neither too short nor too long for a typical bike trip. We introduce a function $g_{u,v}$, based on the Gaussian distribution, defined as follows:

$$g_{u,v} = \exp\left(-\frac{(d_{u,v} - \mu)^2}{2\sigma^2}\right), \quad (2.2)$$

where $d_{u,v}$ represents the distance between stations u and v , with μ and σ denoting the mean and

standard deviation of the network's distance distribution, respectively. The interactions between bike-sharing stations extend beyond binary relationships, necessitating an approach to model the continuum of influence between them. We customized the model to address this complexity by incorporating station-specific features and spatial distances, recognizing the variable impact of proximity on station interactions. This is particularly relevant in scenarios where stations are closely positioned, which might lead to decreased ridership due to the feasibility of walking or opting for other transportation modes for short distances. Conversely, distant stations may also experience a decline in ridership, attributed to the increased physical exertion required. The Gaussian function maps the distances between stations into the model, enhancing its ability to capture the network dynamics and allowing for a customized node embedding formulation tailored to the unique characteristics of the bike-sharing network. By combining GraphSAGE with this distance-sensitive weighting, we leverage the benefits of an inductive framework with an adaptable mechanism that reflects the spatial characteristics of the network.

Figure 2.1 shows the model architecture and framework overview, dividing the methodology into two primary components. The first step is node embedding, which integrates a Gaussian-based distance function with the GraphSAGE algorithm to capture the relationships and feature similarities between stations. The second step is a link weight regression task, aiming to predict ridership between stations based on the derived node embeddings and contextual factors. To commence, we extract various station features and define an initial C -dimensional feature vector for any node z as below:

$$h_z^0 = [x_{z1}, x_{z2}, \dots, x_{zC}], \quad (2.3)$$

where h_z^0 denotes the initial feature vector for node z . Each element x_{zi} represents a unique feature of the station. For example, x_{z1} could represent the number of docks at the station, x_{z2} the number of nearby transit stops, and so on, up to x_{zC} , where C is the total number of features considered.




Figure 2.1: Architecture of the GNN model

We use a deep architecture for node embedding by comprising a series of customized GraphSAGE layers with a depth of d , enabling the representation of learned embedding for each node in the graph to its d -hop neighbors. This approach allows us to enhance the expressiveness of

the learned embeddings. Let $N_s(v)$ denote the set of sampled neighboring nodes of target node v , and h_v^d represent the node v 's embedding at level d . The embedding h_v^d is obtained by incorporating feature messages from neighboring nodes (z), aggregating them based on Gaussian $g_{v,z}$ function, and then combining them with the features of the node v from the previous layer ($d - 1$). This two-stage aggregation can be expressed as follows:

$$h_v^d = \sigma_f(\omega^d \cdot \text{CONCAT}(h_v^{(d-1)}, \text{AGG}(g_{v,z} \cdot h_z^{(d-1)}, z \in N_s(v)))), \quad (2.4)$$

where σ_f is an activation function such as sigmoid or ReLU and ω^d is a learnable weight of d^{th} layer. Subsequently, the features of origin and destination stations at level d are used as an input layer to obtain the link embeddings. Multiple hidden layers using linear and non-linear activation functions are then used to extract progressively weight-relevant information. Finally, an output layer using an activation function predicts the link weight as a continuous number in a supervised prediction setting. We express the predicted link weight for each pair of origin and destination of a link (u and v) as:

$$\hat{y}_{uv} = \sigma_f(\text{CONCAT}(h_u^d, h_v^d)), \quad (2.5)$$

where h_u^d and h_v^d denote the embeddings of the destination and origin stations, respectively. The CONCAT operation combines these node embeddings into a single feature vector, and the activation function introduces non-linearity into the prediction process. This equation enables the model to make continuous predictions using a regression-based approach.

We use data from the peak month of two consecutive years and develop their network (denoted as G_t and G_{t+1}) for training, validating, and testing. Data corresponding to G_{t+1} is specifically used to assess the model's performance with respect to newly introduced nodes and their associated links during network expansion. The model is trained to minimize the discrepancy between predicted and actual bike-sharing station flows as below:

$$\text{Min} \sum_{(u,v) \in E_{w_{uv} > 0}} L(y_{uv}, \hat{y}_{uv}), \quad (2.6)$$

where L is the loss function, u and v are the origin and destination stations of the links, y_{uv}

is the actual flow between two stations, and $\hat{y}_{uv}$ is the predicted flow which is the output of last prediction layer. The specific architecture (i.e., number of layers, activation functions) and hyper-parameters of the model were obtained through a grid search approach while minimizing the loss in the regression problem in a validation data set.

In regression tasks, the Mean Squared Error (MSE) is commonly used to minimize the loss function. Although MSE is a reasonable measure in that it gives the average squared difference between the predicted values and the actual value, it may not always be accurate enough to capture the peak values in data, which is the case in many transportation planning use cases. In the context of bike-sharing, peak demand predictions are considered to satisfy minimum levels of service requirements. We use the Weighted Mean Squared Error (WMSE) as the alternative to MSE to minimize the loss function. WMSE is a variant of MSE commonly used in statistical and Machine Learning literature for regression problems where errors carry different costs or implications. WMSE can be designed to give more weight to errors at higher actual values than lower ones. In this study, we use WMSE to align with the objective of predicting peak demand periods in the bike-sharing network. The WMSE is defined as:

$$\text{WMSE} = \frac{1}{n} \sum_{i=1}^n \eta_i (y_i - \hat{y}_i)^2, \quad (2.7)$$

where η_i denotes the weighting function for the i^{th} data point. The model's specific implementation of WMSE incorporates a linear weighting scheme proportional to the actual flow value y_{uv} , which implies that the weight assigned to each error is directly correlated with the magnitude of the actual flow between stations. Such a weighting approach is instrumental in addressing the challenges of bike-sharing flow prediction, where understanding and predicting high-demand scenarios are paramount for efficient management and user satisfaction. We defined weighting function η_i as:

$$\eta_i = \frac{y_i}{\max(y)}, \quad (2.8)$$

where y_i is the actual value for the i -th data point and $\max(y)$ is the maximum actual value in the dataset. Consequently, we use WMSE to train the model, ensuring that the loss function is min-

imized to accentuate the accuracy of peak demand predictions, thereby tailoring the approach to the pressing needs of the bike-sharing network.

2.4 Data

This section outlines the development of the proposed GNN-WMSE model applied to the Toronto bike-sharing network dataset. We introduce the dataset and the network structure, which are fundamental to the proposed analysis. We also discuss the inherent properties of the network that are expected to affect the model’s performance. We extract station features and relevant network properties, used for developing a comprehensive network representation within the GNN framework. Upon establishing this data structure, we focus on integrating and optimizing the features into the GNN model to enhance its predictive capability.

2.4.1 Data set and Pre-processing

We use the Toronto bike-sharing network as a case study to validate and evaluate the model’s performance. The data was sourced from the Toronto open data portal (<https://open.toronto.ca/>), which provided comprehensive information, including the trip ID, start time, end time, trip duration, and origin and destination station name. The dataset covers the period from 2016 to 2022, encompassing monthly data. Before applying the model, we performed descriptive data analysis to explore trends, identify anomalies, and assess the completeness of the data. Missing values in station information were sourced from complementary resources, while missing values in trip flows were few in number and removed to avoid errors. Outliers, such as unusually short or long trip duration, were identified and excluded to prevent distortion in the analysis. During the pre-processing steps, the dataset was cleaned, completed, and properly structured to the model’s requirements.

Figure 2.2 represents the distribution of Toronto’s bike-sharing stations, categorized by their respective years of development. The figure offers insights into the bike-sharing network’s growth and expansion, showcasing the locations of stations across the city. The bike-sharing network in Toronto has consistently increased in popularity and usage over the past seven years. The rise

in total ridership substantiates this growth, as depicted in Figure 2.3-a. The figure illustrates the consistent upward trajectory of total ridership, with an average annual increase of 27%. Analyzing the monthly trend of bike-sharing trips, as shown in Figure 2.3-b, we observe that August consistently records the highest flow, followed by July, June, and September. This finding confirms the correlation between bike-sharing demand and weather patterns, with warmer months experiencing more trips. Consequently, the analysis focuses on a monthly prediction strategy specifically targeting the peak month of August.




Figure 2.2: Bike-sharing growth map in Toronto

To determine the appropriate distance threshold for constructing the adjacency matrix of the bike-sharing network, we analyze the distance histogram and cumulative distribution function (CDF), as illustrated in Figure 2.3-c. We observe that more than 90% of the data points fell within a distance of 12 km. Therefore, we select this distance as the threshold value, which captures a significant portion of the data, signifying the presence of a link between the two stations in the graph. Using this threshold, we construct the adjacency matrix (A) that is the basis for capturing the connections between bike-sharing stations. Moreover, this distance threshold aligns with the 30-minute time limit imposed by the Toronto bike-sharing system, where a typical cycling speed of 20–25 km/h allows users to cover approximately 12 km within the permitted time. By incorporating both data-driven insights and operational constraints, this

adjacency matrix definition ensures a realistic and practical representation of the bike-sharing network.




Figure 2.3: Toronto 2016-2022, a)ridership, b)monthly trips, c)distance cumulative distribution function (CDF)

We analyze the flow between station pairs during the peak month, identifying that some pairs have minimal trip utilization. To ensure the focus of the model is on the most significant flow patterns, we exclude station pairs with fewer than five trips per month. This decision was guided by an analysis of the trip distribution, which demonstrated that these low-utilization pairs contributed negligibly to the overall network dynamics. While our approach focuses on removing underutilized station pairs, Lin et al. (2018) adopted a similar rationale by excluding rarely utilized stations from their analysis of New York City’s bike-sharing network to ensure the robustness of their demand prediction model. By filtering these pairs, the model can better represent the dominant and meaningful flows within the bike-sharing system. Consequently, we refine the analysis by updating the flow matrix based on the constructed adjacency matrix. In this refined flow matrix, denoted as W , each element (w_{uv}) represents the corresponding flow between two stations (u and v) that are directly connected by a link within the network. Note that the flow from station u to station v (w_{uv}) is distinct from the flow in the reverse direction (w_{vu}), as this study aims to capture the directed nature of the bike-sharing system.

Additionally, by examining the relationship between trip frequency and distance, we discern a range wherein most trips predominantly occur. We show that stations located at intermediate distances from each other are more frequently traveled, while those that are either excessively close or far tend to have lower trip frequencies. This observation corroborates the Gaussian approach adopted for customizing the node embedding process. With its peak representing the most common trip distances, the Gaussian function effectively captures this trend in the data. Consequently, we apply this knowledge to fine-tune the model by fitting a Gaussian distribution to the trip data, allowing us to determine the optimal values for the parameters sigma σ and mu μ . These parameters are instrumental in refining the node embeddings to mirror the real-world trip distribution, further enhancing the model’s predictive accuracy in reflecting the operational patterns of the Toronto bike-sharing system.

2.4.2 Station Features

Bike-sharing stations possess attributes that provide insights from various perspectives. These attributes encompass the number of stalls and the availability of e-bikes within each station. Demographic sets of information such as population, age, and income are sourced from ward profiles and Census data to enhance the dataset. Furthermore, using the derived flow matrix in the previous section, we introduce a super-station indicator that aggregates the total number of inbound and outbound flows, enabling the identification of the crowded stations based on flow value. Moreover, this study considers POIs surrounding each station and incorporates station (nodes) characteristics derived from the network analysis, such as nodal strength. Additionally, we account for the spatial relationship between stations by incorporating the distance between them. This feature is integrated into the model through a Gaussian-based weighting function, enabling the capture of inter-station dynamics and their influence on ridership patterns. In the following sections, we delve into a detailed discussion of these two features and their implications within the study.

POIs is a designated location containing a range of establishments, such as tourist attractions and transportation hubs. The existing literature highlights the proximity of POIs as influential factors that contribute to travel patterns in bike-sharing networks (Liu et al., 2017a). Conse-

quently, we incorporate POIs as a feature of nodes in the model. In this study, we acquire POI data utilizing the MapBox Tilequery API (map, 2020) using transit stops and POI labels, see Liu et al. (2025) and Chiou and Wu (2024) for transit-integrated bike-sharing. Specifically, we consider a circular region with a radius centered at each bike station. To determine an appropriate value for R , we refer to relevant literature (Xiao et al., 2023) and consider an acceptable average walking distance. Consequently, we select a radius of 2.5 KM as the most representative measure of the station’s surrounding area. Subsequently, we count the number of POIs within each station’s circular boundary and categorize them. Table 2.1 demonstrates the type of POIs considered for each category. As part of the node features, we use a vector that expresses the count of POIs within each station’s vicinity for each specific category.

Table 2.1: POIs Categories

Category	POIs type examples
Financial	Banks, Offices (i.e., insurance office, law office, real estate, and accountant office)
Commercial	Shops, Stores, Malls, Marketplace, Commercial Area, Cafe, Restaurants, Bars
Education	Schools, Driving School, Library, College, University, Educational Institution
Recreational	Sports Complex, Playground, Tourist Attraction, Beach, Theater, Cinema, Zoo
Residential	Residential Buildings and Areas
Healthcare	Doctor’s Office, Pharmacy, Dentist, Hospital, Clinic, Drugstore
Transportation	Parking, Charging Stations, Harbors, Transit Stops
Other	Fire Station, Post Office, Police, Mosque, Synagogue, Church, Guidepost

The characteristics of the bike-sharing network provide deeper insights into its functionality and can be used to describe various aspects of the network, including its connectivity and centrality. As a result, we extract the station network properties and assign them as part of the features. Table 2.2 provides an overview of the network’s characteristics in two consecutive years, including the size, topological properties, evolutionary trends, and essential measures such as the average degree (K), average shortest path length (d), network diameter (D), triangle count (T), and clustering coefficient (C_s). The shortest path length represents the minimum number of links required to travel between two nodes, and the network diameter refers to the longest

shortest path. Triangle count indicates the local clustering by counting the closed loops of three interconnected nodes, and the clustering coefficient measures the degree to which nodes tend to cluster together- see Newman (2018) for detailed information about network analysis techniques. The values presented in this table are calculated while considering the network as a weighted directed network.

Table 2.2: Bike-sharing Network Properties

Network	$ V $	$ E_{w_{ij}>0} $	K	d	D	C_s	T
G_{2022}	654	40035	61.22	2.08	9	0.48	389587
G_{2021}	589	25494	43.28	2.44	11	0.36	169717

Based on the network properties, the bike-sharing network exhibits considerable growth in nodes and usage and forms a dense graph structure with a relatively high density of connections among the stations, especially in G_{2022} . The observed value of the clustering coefficient and triangle counts coupled with the small shortest path further admit the network’s dense interconnectivity pattern and dense nature.

Following these network properties, we capture four additional measures that provide valuable insights into the importance, connectivity patterns, and influence of nodes within the bike-sharing network. First, nodal strength is the measure of a node’s influence, which reflects the overall impact by including the weights of its links. Second, PageRank identifies influential nodes based on the quality and quantity of incoming links. Third, the weighted clustering coefficient assesses the tendency of a node’s neighbors to form clustered connections. Lastly, betweenness centrality highlights a node’s significance as a bridge or bottleneck within the network by measuring the extent to which a node lies on the shortest paths between other pairs of nodes.

Figure 2.4 displays the variations of the aforementioned station measures. The difference between the median and average nodal strength shows an uneven distribution of influence among the stations with few highly influential ones (Figure 2.4-a). Furthermore, some stations are closely connected to their neighbors, forming tightly clustered connections, as shown in Figure 2.4-b. Lastly, Figure 2.4-c highlights that some stations act as critical bridges with high betweenness centrality. Consequently, we incorporate these measures into the station features in the model.




Figure 2.4: Stations' Measure Variations in the Network, a)Nodal strength, b)Weighted clustering coefficient, c)Betweenness centrality

2.5 Analysis

We use the data for the peak month in years 2021 (G_{2021}) for training the model and consider the new nodes and their associated links in the same month of year 2022 (G_{2022}) as the network expansion, serving as the unseen dataset for evaluating the final model's performance. The experiment is conducted using Python 3.0, and the models are developed using TensorFlow, an open-source library designed for high-performance numerical computation, particularly well-suited for deep learning models. We tune the hyperparameters, which are the settings and configurations that govern the model's learning process, to balance model complexity and optimal performance as part of the training process. The grid search systematically explored all combinations of hyperparameters within the specified ranges. Each configuration was evaluated using a validation dataset, with performance assessed based on validation accuracy and loss. The optimal values reported in Table 2.3 correspond to the settings that achieved the best performance across these metrics, ensuring robust and generalizable results. This table summarizes the hyperparameter grid search range and the obtained optimal values, which are used in

the model. The first four hyperparameters shape the GNN model structure and capacity. The following three hyperparameters determine the optimization algorithm and learning process, while the final three are designed to prevent overfitting. Following hyperparameter tuning, we retain the node features that contribute to enhancing the model’s performance.

Table 2.3: Hyperparameter Tuning

Num.	Hyperparameter	Optimal Value	Grid Search Space
1	Number of layers	6	(1,2,3,...,10)
2	neighborhood search depth	10	(1,2,3,...,D)
3	Node embedding size	16 (in), 32 (out)	(10,...,60)
4	Activation function	ReLU	(ReLU, Linear, Tanh, Sigmoid)
5	Epochs	300	(1-500)
6	Optimizer	Adam	(Adam, SGD, RMSProp)
7	Learning rate	0.001	(.000001, .00001, .0001, .001)
8	Dropout	0	(0,0.1,0.2)
9	Weight Decay	0.001	(0.1,0.01,0.001,0.0001)
10	L2 Regularization	0.1	(0.001,0.01,0.1,0)

Our goal is to minimize the discrepancy between the predicted values and the ground truth labels. To evaluate the efficiency of the model, we use two primary metrics: Root Mean Square Error (RMSE) and Mean Absolute Error (MAE). These metrics are compared with those of benchmark models, which will be explained in the next section. The definitions of these metrics are as follows:

$$RMSE = \sqrt{\frac{1}{N} \sum_{i=1}^N (y_i - \hat{y}_i)^2}, \quad (2.9)$$

$$MAE = \frac{1}{N} \sum_{i=1}^N |y_i - \hat{y}_i|, \quad (2.10)$$

where N represents the total instances, y_i denotes the actual flow of each trip (link), and $\hat{y}_i$ represents the predicted flow.

2.5.1 Benchmarks

Previous studies on demand forecasting have primarily used statistical regression methods or tree-based algorithms in Machine Learning. Additionally, several deep learning architectures were explored in the existing literature. We conducted a benchmarking exercise to evaluate the

effectiveness of the proposed GNN-WMSE model and calculate RMSE and MAE for all models to compare them. This provides a quantitative basis to assess the performance of models in terms of prediction accuracy and reliability. We used the same dataset division across all models to ensure consistency and comparability. Since the goal is to address long-term bike-sharing network planning rather than short-term demand fluctuations or time-series dependencies, we compare the proposed model to the following established baseline models:

1-Linear Regression (LR): This model uses a regression approach, acknowledging a simple linear relationship between input features of bike-sharing stations and their transactions.

2-Spatial Regression (SLR): This model uses a regression approach, acknowledging spatial relationships in addition to linear regression model. Spatial Regression extends the capabilities of Linear Regression by integrating spatial dependencies, which is essential for capturing patterns in bike-sharing data.

2-XGBoost: This tree-based Machine Learning model leverages the gradient-boosting decision tree mechanism. It utilizes the same input parameters as the Linear Regression model, providing a robust comparative framework. XGBoost introduces a non-linear approach that can capture complex interactions between features, serving as a robust benchmark against which the predictive power of GNN-WMSE model can be measured.

3-Artificial Neural Network (ANN): To examine the influence of incorporating d-hop station neighbor features on predicting the flow between two bike-sharing stations, we use an ANN algorithm that does not take these neighbor features into account. This model evaluates the performance of a traditional non-graph-based approach, allowing us to contrast its effectiveness with the graph-based nature of the proposed model in the bike-sharing network. In this model, each link in the data set is represented by a pair of nodes. The architecture of a simple ANN model consists of an input layer, several hidden layers, and an output layer. As illustrated in Figure 2.5, we concatenate the features of the origin and destination stations to determine the link features as a vector of size $2C$, where C represents the dimension of the individual station features. This concatenated link feature vector serves as the input layer for the ANN model. The model is trained in a supervised manner, where it learns to map the input (link features) to the output (link weights), leveraging the available training data.




Figure 2.5: Architecture of the ANN model

Table 2.4: Models performance comparison for G_{2021} and G_{2022}

Models	Existing G_{2021}		New Stations G_{2022}	
	RMSE	MAE	RMSE	MAE
LR	18.93	8.34	26.16	8.67
SLR	18.15	8.29	25.67	11.07
XGBoost	16.56	8.01	24.56	8.35
ANN	15.47	8.03	24.12	8.05
GNN-WMSE	<u>14.03</u>	<u>7.12</u>	<u>21.85</u>	<u>7.43</u>

The results of different models for the existing network G_{2021} are presented in the second column of Table 2.4. Compared with the baseline models, the proposed approach outperforms the baseline models in terms of the key evaluation metrics. The GNN-WMSE model has improved the RMSE value by an average of 18% over all of the benchmarks and approximately 15% compared to the XGBoost model, which is the best-performing model among the non-deep learning established benchmarks. This advancement underscores the ability of the proposed model to handle complex data more effectively than conventional Machine Learning approaches. Deep learning techniques can capture complex, non-linear relationships. Furthermore, incorporating a graph structure in the network model allows for a comprehensive capture of features related to $d - hop$ stations. Another distinctive advantage of our model is the integration of GraphSAGE, which has been further customized to account for spatial interactions within the learning process. This addition significantly enriches the model's ability to understand and predict the dynamics of the network more accurately. Regression models, including the Spatial Linear

Regression (SLR) model, displayed inferior performance compared to other baselines. This outcome can largely be attributed to their oversimplified assumptions, which struggle to handle the intricate complexities and interdependencies present in the dataset. The shortcomings of these traditional models in dealing with non-linear and spatially complex data highlight the necessity for more advanced approaches. In comparison, the XGBoost model demonstrated a slight improvement in prediction accuracy over the linear regression models, as evidenced by both the MAE and RMSE metrics. This improvement indicates the effectiveness of tree-based methods in capturing more complex relationships than traditional linear models. The results show that the inclusion of deep learning techniques (ANN model) enhanced prediction performance. This model, which does not incorporate a graph structure, nonetheless slightly better captured complex nonlinear relationships within the bike-sharing network.

The comparison of the ANN and GNN-WMSE models shows that the graph-based model (GNN-WMSE) offers superior predictive power, as indicated by its lower RMSE and MAE values. Figure 2.6 illustrates the comparative performance of these two models across a subset of two hundred links, randomly selected for visual representation purposes. This advantage is attributed to the GNN model's ability to incorporate neighborhood information through node embedding, a key factor in predicting bike-sharing flows. Unlike the ANN model, which processes data in a tabular format, the GNN-WMSE model effectively utilizes the graph structure of the data, learning station representations that account for both individual features and their $d - hop$ station neighbors simultaneously.




Figure 2.6: Predicted Versus Actual Flows in Neural Network Models: a) ANN (Artificial Neural Network), b) GNN (Graph Neural Network).

Upon evaluating the results from the benchmarks, a significant concern emerged with some of the predictive approaches, notably the LR, SLR, and XGBoost models. These models displayed a tendency to forecast negative values for flows in bike-sharing systems, which are intrinsically non-negative. Although these negative predictions are statistically possible within the models, they are practically infeasible and do not align with the actual nature of the data. Figure 2.7 illustrates this discrepancy, presenting a distribution plot of the models' predictions. The red area in the chart shows the frequency and extent of the negative predictions made by the models, indicating that these models are unsuitable for the given data. This visualization emphasizes the need for a predictive approach that guarantees all forecasts fall within the plausible range of non-negative values.




Figure 2.7: Incidence of Negative Predictions in LR, SLR, and XGBoost Models

Figure 2.8 compares the predictions of models with the actual trip numbers for a subset of the dataset. These models use different methodologies, such as basic regression, advanced tree-based, and graph neural network techniques. This comparison captures performance metrics, such as the overestimation or underestimation of predictions and the variance in model performance. The figure reveals that the GNN-WMSE Model has a tighter clustering of predictions around the identity line for most data points, especially in the mid to high range of actual flows. This indicates a strong correlation between the predicted and actual values. In instances of lower-valued flows, the GNN-WMSE Model tends to underestimate their magnitude. Nonetheless, it exhibits superior prediction accuracy overall, as evidenced by its closer adherence to the

identity line than the other models. The XGBoost model displays a dispersion of points, especially in the lower range of true values, where it significantly overestimates trip numbers. The Linear Regression model performed the worst among the three models. Therefore, the GNN-WMSE model provides the most reliable predictions, striking an effective balance between precision and generalization.




Figure 2.8: The relationship between true and predicted flow

We also focused on comparing the performance of the GNN-WMSE model with the XGBoost model in specifically predicting high flow values, the accuracy of which is essential to managing systems operations and enhancing bike-sharing availability. Figure 2.9 shows the comparative performance of these two models under conditions of peak flows. The data is sorted from high to low based on actual flow values. Figure 2.9-a displays a trend of predictions with respect to true flow, with the horizontal axis representing the index of observations sorted by actual flows and the vertical axis showing smoothed ratios of actual to predicted flows. We used a rolling average and data smoothing technique to enhance the clarity of trends. This approach ensures a more balanced view of the dataset by averaging over a specified window, aiding in the reliable detection of performance trends across the dataset. In this visualization, the XGBoost

model's trend line, positioned above that of the GNN-WMSE, suggests it more frequently or more significantly underestimates the flow values compared to its counterpart. This tendency is especially pronounced at higher flow values, which are of primary interest in the analysis. We shaded the region between the two trend lines for easy comparison of predictive performance. The performance gap decreases as we move to lower actual flow values, but the distinction at higher values, although slight, is in line with the target objectives for scenarios of peak flow demand. Additionally, Figure 2.9-b presents this comparison differently. Here, the x-axis shows the actual flow values sorted from high to low, while the y-axis represents the difference in performance between the XGBoost and GNN-WMSE models. The red points mean that the XGBoost model has a more significant error than the GNN-WMSE model in the corresponding instances. This figure reveals that the deviation in the GNN-WMSE's performance is generally lower than XGBoost, which makes the GNN-WMSE model more stable across true values. The trend line demonstrates an increasing discrepancy in predictive performance as the flow values rise, emphasizing the GNN-WMSE model's superior precision and consistency, particularly in situations of critical peak demand.





Figure 2.9: a)Comparative Performance Trends of GNN-WMSE and XGBoost (best-performed benchmark model) in Predicting High Flow Values; b)Differential Performance Analysis of GNN-WMSE and XGBoost (best-performed benchmark model) in Predicting High Flow Values

Note that the prediction of bike-sharing flow using the GNN-WMSE model still presents opportunities for improvement. Several factors could explain this limited predictive power. For instance, as discussed, the bike-sharing network is highly dense, with over twenty-five thousand links and only about six hundred nodes. In such dense networks, the complexity of the relationships increases and makes it challenging to understand the pattern. Furthermore, such dense networks are prone to overfitting, where the model becomes excessively complex and tends to

memorize the training data rather than capture general patterns. These factors directly impact the accuracy of predictions for unseen and new data. In the following section, we evaluate the performance of the model in network expansion.

2.5.2 Link Weight Prediction for G_{2022}

We evaluate the model’s applicability in predicting flow patterns during network expansion by leveraging the trained model to forecast bike-sharing flows for newly introduced nodes and their associated links. To this end, we assess the performance of the trained model using the G_{2022} network, focusing on the same peak month in which 65 new nodes have been added. Using the trained GNN-WMSE model G_{2021} , we predict the weights of the new links, which represent the bike-sharing flows between the newly introduced stations and the existing ones. We note that the model assumes that bike-sharing users do not change their routes by hopping between stations, as bike-sharing systems usually impose a 30-minute usage limit per trip, indicating that users are unlikely to switch bikes at different stations during a single journey. This assumption simplifies the network expansion prediction process and is essential for interpreting the model’s flow predictions in the context of real bike-sharing usage patterns.

The evaluation and benchmarking outcomes for the network expansion in G_{2022} are presented in the third column of Table 2.4. The results indicate that the GNN-WMSE model achieves an average improvement of 15% in RMSE and 20% in MAE over all of the benchmarks. As demonstrated, the GNN-WMSE model consistently exhibits lower RMSE and MAE values, reflecting a comparable accuracy in predicting flows for both the existing links in G_{2021} and the newly established links in G_{2022} . This continuity of performance underscores the model’s robustness across network expansion. Additionally, Figure 2.10 visualizes the comparison, highlighting the variance in the top-500 precise link weight predictions between the two graphs, G_{2021} and G_{2022} .




Figure 2.10: Top-500 accurate link weight predictions versus actual weights- a)2022, b)2021.

The analysis yields several insightful observations. Firstly, links associated with new nodes displayed a higher accuracy in predicting bike-sharing flows, particularly when these nodes exhibited high nodal strength. Secondly, accuracy was further enhanced for new bike-sharing stations located near transportation and transit Points of POIs and those with higher stall capacities. These factors collectively played a significant role in improving the accuracy of flow predictions for new links.

A comparative examination of the top-500 accurate link weight predictions in G_{2021} revealed a similar trend: nodes with above-average stall capacities and high nodal strength values were more likely to yield accurate link weight predictions. Furthermore, these predictions were most accurate for links connected to nodes in medium flow ranges, with fewer predictions pertaining to high-volume trips. This pattern underscores the importance of integrating network analysis outcomes, transportation infrastructure characteristics, and stall capacity into the predictive modeling process. These findings validate our approach and highlight aspects that could guide future network planning and expansion strategies.

2.5.3 Expansion Design

To bridge the gap between predictive modeling and practical decision-making, we developed an interactive decision-support tool designed to guide network expansion planning. While the full design and technical implementation of the tool are described in Chapter 5, this section high-

lights its role in supporting expansion scenario analysis grounded in the GNN-WMSE ridership prediction model.

The tool enables the extraction of candidate locations and relevant station features for proposed expansions, leveraging the same data sources and processing pipelines used in the modeling components of this chapter. By integrating ridership predictions, demographic characteristics, and transit accessibility indicators, the tool supports scenario-based analysis for strategic planning.

We focus on determining the optimal locations for a new bike share station within the G_{2022} network with an average of 15 stalls. We aim to pinpoint the location with the highest potential for increasing bike-sharing ridership using the tool. By leveraging the GNN-WMSE model, we predict the bike-sharing flow between the new node and existing ones and rank the possible locations. Five distinct scenarios for determining the location are considered to capture the essence of each location type and link directly to policy actions that can be pursued to maximize the efficiency of bike-sharing systems. Urban Core Connectivity targets densely populated areas with advanced transportation networks that facilitate seamless intermodal transfers, aiming to reduce congestion and improve commute efficiency. Suburban Shortfalls focuses on rapidly growing suburbs that struggle with inadequate public transport options, leading to increased reliance on private vehicles and higher traffic congestion; introducing sustainable transit options like bike-sharing in these areas with low transportation infrastructure can bridge these gaps and reduce environmental impact. Transit-Enhanced Suburbs emphasizes emerging residential zones where recent transportation infrastructure investments have not been matched by population density, enhancing connectivity for commuters who rely on public transportation and fostering community growth and business attraction. Equity Expansion seeks to promote equitable access to bike-sharing services in underserved areas with low transportation infrastructure, often marginalized and economically disadvantaged neighborhoods where mobility and access to essential services are limited. Lastly, the Tourist-focused design concentrates on recreational and tourist destinations with high-value points of interest, catering to the transportation needs of visitors and supporting vibrant centers of activity that attract tourists and locals to cultural, historical, and recreational attractions.

Figure 2.11 demonstrates the flows coming in and out of the new station under the aforementioned five scenarios. The results indicate that selecting the location based on the Urban Core Connectivity policy would generate a substantial volume of trip initiations from this new station and result in a 6% increase in total ridership. Conversely, Equity Expansion scenario yields a less favorable outcome, with a 3.38% increase in ridership. While this scenario focuses on providing equitable access to underserved areas, it experiences lower ridership flows. This indicates that although it successfully addresses transportation gaps, the overall increase in ridership is modest compared to more densely populated or well-connected regions. The results support the decision to prioritize locations like those in the Urban Core Connectivity scenario for new bike share stations, as these areas foster higher usage rates and contribute substantially to the overall network ridership growth, likely due to dense populations and well-developed transportation infrastructure.

Both the Suburban Shortfalls and Transit-Enhanced Suburbs scenarios show moderate flow levels. While the Suburban Shortfalls scenario features less developed transportation infrastructure compared to the Urban Core, it still attracts significant incoming flows, suggesting effective utilization of bike-sharing as an alternative transport mode. The Transit-Enhanced Suburbs scenario, despite its low population density, maintains balanced incoming and outgoing flows, likely due to strong connectivity features that attract commuters. Lastly, the Tourist-focused scenario exhibits a significant disparity between incoming and outgoing flows. This substantial bi-directional flow suggests that the location attracts tourists and serves as a starting point for their local travels, making it a dynamic node within the bike-sharing network.




Figure 2.11: Expansion scenarios results

2.6 Conclusions

As bike-sharing continues to grow and expand, it becomes increasingly important to analyze and accurately model the spatial distribution of demand and ridership across the network. Current bike-sharing ridership prediction models focus on aggregated ridership inflow and outflow, missing the ability to provide station-to-station trip distributions. This study uses a GNN to predict ridership between bike-sharing stations, capturing the impacts of network-level features. The GNN provides a link weight prediction framework in a weighted directed network. We integrate neighborhood information through node embedding by leveraging GraphSAGE and distance Gaussian-based functions. Consequently, link weight regression prediction is conducted using the generated link embeddings corresponding to the origin and destination stations.

Using Bikeshare Toronto as a case study, we present a supervised learning framework to assess ridership patterns, contrasting the GNN-WMSE model against four benchmarks: Linear Regression, Spatial Regression, XGBoost, and Artificial Neural Network. The GNN-WMSE model, which integrates neighborhood information via customized node embedding using GraphSAGE, emerges as a superior method for predicting bike-sharing flows. This model's proficiency emphasizes embedding graph-based features to effectively capture the complexities of

bike-sharing usage patterns.

The prediction capability of the GNN and benchmarks was further examined in network expansion scenarios, where the models were tasked with forecasting ridership of newly added stations during the same peak month in the subsequent year. Additionally, the model's performance in predicting high flow values is evaluated, recognizing that such accuracy is important for managing system operations and improving bike-sharing availability. The results indicate the GNN-WMSE model's superior precision and consistency, with an average improvement of 18% in RMSE and 12% in MAE for existing stations. For new stations in the following year, it achieves an average improvement of 15% in RMSE and 20% in MAE. A visual analysis of peak demand predictions suggests that the GNN-WMSE model achieves a notable reduction in error deviations for the highest flow values compared to the next-best model, XGBoost, underscoring its reliability in such scenarios. Moreover, the model supports scenario analysis to identify optimal locations for new stations to enhance overall ridership.

The bike-sharing data used in this study is publicly available for most cities, with station and ridership information accessible through the General Bikeshare Feed Specification (GBFS) data format and various open data portals. Furthermore, supplementary information such as census data and POIs can be obtained from open data portals and APIs. Bike-sharing data availability improves the model's scalability to other cities, reinforcing its relevance and adaptability beyond the specific case study of Toronto and Vancouver. While the proposed model effectively predicts bike-sharing ridership and informs network expansion, several limitations should be noted, including the exclusion of low-flow patterns, the focus on peak periods, and the absence of behavioral and temporal factors. Future research should address these limitations by integrating optimization algorithms for selecting n -optimal new station locations to further enhance prediction accuracy and inform strategic bike-sharing network planning. This integrated approach will offer a data-driven framework for designing and expanding future bike-sharing networks.

Chapter 3

Station Placement Optimization

Abstract: Bike-sharing is an active transportation mode that can be used independently or jointly with other modes to offer first- and last-mile connectivity. Optimizing bike-sharing network expansion is complex, particularly when considering ridership maximization as an objective, which introduces network-level effects in the decision-making process. Discrete optimization methods become impractical for large-scale problems with numerous potential station locations, especially when detailed data on location properties is lacking. This study introduces a Continuum Approximation (CA) model with a force-based algorithm to configure the placement of bike-sharing stations that maximize the ridership. Unlike previous approaches focused primarily on accessibility, our model integrates ridership objectives into the CA framework, allowing stations to move freely across a continuous service area where origins and destinations are spatially distributed. The proposed algorithm extends the classic force-based facility location model by incorporating both repulsion and attraction forces to optimize station spacing and dynamically align station locations with high-demand areas. Numerical experiments demonstrate the effectiveness of the CA algorithm, delivering near-optimal solutions that enhance bike-sharing ridership. Sensitivity analysis further validates the adaptability of the model, revealing that optimizing station spacing and balancing interactions between inter-station and demand-to-station dynamics are crucial for overall system performance. The results identify an optimal parameter range that maximizes ridership while ensuring the model's responsiveness to varying demand densities and network sizes. Subsequently, to illustrate the practical applicability of this study approach, we apply the proposed CA model and algorithm to the real-world context of the Vancouver bike-sharing system.

Keywords: Facility Location Problem, Continuum Approximation, Force-based Algorithm, Bike-Sharing, Ridership Optimization.

3.1 Introduction

Bike-sharing is an active mode of transportation used either independently or jointly with other modes like transit to provide first- and last-mile connectivity. Bike-sharing services have been expanding in many cities in response to growing demand as more population segments become environmentally conscious and tend toward affordable non-ownership mobility options. Systematic expansion of a bike-sharing network can be challenging, as it depends on the current state of the network, land use and socio-demographic features of the city, and the type of objective(s) considered for evaluating expansion plans. The four primary bike-sharing network expansion objectives include accessibility, transit integration, equity, and ridership promotion, frequently used in bike-sharing master plans commissioned every 3-5 years to plan the network's expansion. Among the indicators, the first three are location-specific, whereas the fourth, riding promotion, is network-specific, as it depends on the strategic location of new stations with respect to the rest of the network.

Traditional bike-sharing network design problems focus on the location-specific objectives (Frade and Ribeiro, 2015) and often overlook the interaction between stations in promoting ridership. Wang et al. (2016) used spatiotemporal analysis of historical bike-sharing trip data to identify demand patterns in station utilization, locate areas with a shortage of bikes and/or bike racks, and find optimal locations for new stations. Some studies, such as Larsen et al. (2013) and Banerjee et al. (2020) integrated demand indicators and used Geographic Information System (GIS)-based approaches to determine the optimal bike-sharing stations, while others, like Martinez et al. (2012), have applied heuristic algorithms to optimize station locations alongside fleet size and relocation logistics. While these studies have demand as an output of the desired design, they do not focus on maximizing ridership directly. Hence, a gap in the bike-sharing network expansion literature is to develop ridership maximizing network design optimization models. Further enhancing the applicability of such a model is to make it as data-driven as possible so the recommended optimal station locations consistently project their expected demand if deployed in reality.

Optimization of a bike-sharing network expansion plan is complex when ridership promotion

is the objective because of the inherent network-level effects. In the literature on facility location problems, Mixed Integer Programming is a widely used methodology to determine station locations from a finite set of candidates (Drezner and Hamacher, 2004). However, discrete optimization methods are impractical for large-scale networks with many potential locations due to the NP-hard nature of the problem. While several studies have applied Mixed Integer Programming to bike-sharing network design, most focus on operational considerations such as cost minimization rather than explicitly incorporating ridership demand dynamics into the optimization framework. This gap highlights a key challenge in strategic- and tactical-level planning, in which the level and quality of available data are not necessary for high-level decision-making, complicating the discrete optimization approach, which would require a reasonably large number of potential locations accompanied by other socio-demographic and land-use information for each location.

Continuum Approximation (CA) was initially proposed to handle some of the implementation challenges of discrete facility location problems (Newell, 1971, 1973; Daganzo, 1984; Daganzo and Newell, 1986). CA simplifies facility location optimization by treating demand and other variables as continuous density functions, allowing analytical solutions to identify optimal facility placements. The strategy of CA is to use smooth density functions of demand and other socio-demographic features to find the optimal locations of a set of facilities that can move freely within the feasible boundaries of a surface. A seminal application of CA is a study of Ouyang and Daganzo (2006), where force-based adjustments are used to achieve near-optimal placements for a set of terminals. The method of Ouyang and Daganzo (2006) considers two types of repulsion forces: one acting along the connection between the centers of overlapping influence areas of terminals (as circular disks) and the other from the boundary of service area towards the terminal influence areas. This methodology automates the adjustment of these disks within the service area until equilibrium is reached to achieve near-optimal locations. In another related study, Ouyang (2007) proposed an algorithm to obtain near-optimum vehicle routing zones by automating the partitioning of delivery zones within the service area. The algorithm incorporates the forces described in (Ouyang and Daganzo, 2006) to adjust the zone geometries toward an optimal configuration. Several other studies have applied CA techniques

to optimize various aspects of bike-sharing systems, such as station placement, fleet size, parking slot allocation, and operational efficiency (Luo et al., 2021; Soriguera and Jiménez-Meroño, 2020; Zhang et al., 2023). While these studies focus primarily on cost minimization and strategic planning, none have combined CA with a force-based algorithm to maximize ridership in bike-sharing networks directly.

The concept of using CA for network design in the presence of two-dimensional origin-destination demand patterns has been previously introduced in applications such as transit network design (Ouyang et al., 2014), where demand-dependent route densities are optimized to reduce costs and improve adaptability. Similarly, Caicedo et al. (2023) focused on designing a cost-effective bike network that accounts for variable travel demands across different areas. CA is also applied to flow models of shared ownership, such as dynamic electric vehicle (EV) sharing, considering fluctuating demand to optimize EV sharing stations' locations (Li et al., 2016b). The common thread among these flow-capturing models of CA is a structured network in the background, commonly a Manhattan grid, that forms a lattice for the emerging network. In hindsight, the models lack flexibility in the network structure, like that achieved in the Ouyang and Daganzo (2006) model, which allows disks to move freely in space based on applied forces.

This study proposes a CA model with a force-based algorithm that maximizes bike-sharing ridership by moving the stations based on the applied forces until equilibrium is reached. Unlike Ouyang and Daganzo (2006), which has only repulsion forces for coverage-maximization facility location, we show that the ridership-maximization facility location problem has both repulsion and attraction forces, with the latter pushing stations to points of high origin and destination (OD) demands. The proposed force-based heuristic algorithm strategically adjusts each station's location within a continuous service area, where demand is spatially distributed, to align with OD patterns and maximize network ridership. The proposed force-based heuristic algorithm strategically optimizes station placement within a continuous service area, allowing stations to move freely and adapt to spatial demand patterns. This methodology places ridership at the forefront, integrating demand-driven optimization into the CA framework, which goes beyond the traditional emphasis on coverage. This approach supports tactical-level planning and improves network efficiency by enabling stations to explore and settle at optimal locations. Nu-

merical experiments demonstrate the effectiveness of the CA algorithm, delivering near-optimal solutions that enhance bike-sharing ridership. The results of the numerical experiments indicate that optimal station positioning hinges on achieving a balance between interactions among stations and between stations and high-demand areas. Further analysis has identified an optimal station spacing and adherence range, significantly boosting ridership. Testing the algorithm's performance in larger-scale scenarios and against various spatial demand patterns has confirmed its adaptability to diverse operational scales and real-world complexities.

The remainder of this chapter is organized as follows. Section 3.2 presents a comprehensive review of the relevant literature, highlighting key methodological advancements and positioning our contribution within the existing body of work. Sections 3.3 and 3.4 describe the problem formulation and the proposed solution methodology, respectively. In Section 3.5, we present the results of a computational experiment, analyze the model's performance, and conduct sensitivity analyses to evaluate its robustness. Section 3.6 illustrates the model's applicability through a real-world case study of the Vancouver bike-sharing system. Finally, Section 3.7 concludes the chapter by summarizing key insights and outlining potential avenues for future research.

3.2 Literature Review

This section reviews the relevant literature on bike-sharing station location optimization, examining diverse strategies for their placement while considering ridership impacts. Subsequently, we provide a comprehensive overview of Continuum Approximation algorithms for facility location problems and their applications in other logistics settings.

3.2.1 Bike-Sharing Station Location Optimization

Strategic placement of bike-sharing stations significantly affects ridership and optimization techniques can be used to enhance this effect. Bike-sharing station location optimization literature covers methodologies such as mathematical algorithms, multi-criteria decision-making, and Geographic Information Systems (GIS)-based algorithms (Bahadori et al., 2021).

Effective budget allocation is important for determining the number and locations of bike-

sharing stations. Frade and Ribeiro (2015) used a maximum coverage model to determine station locations in response to induced demand while considering budget constraints. Qian et al. (2022) focus on equity as the objective of network expansion and propose a multi-objective optimization model that uses a genetic algorithm to balance the distribution of stations, aiming to optimize potential demand, annual revenue, and accessibility to address equity. Similarly, Caggiani et al. (2020) integrated bike-sharing with public transit systems to enhance overall accessibility and multimodal connectivity. The authors developed a model to minimize a statistical measure index representing inequalities in accessibility while simultaneously maintaining specified levels of accessibility and coverage. The results revealed that focusing solely on accessibility or coverage, without considering equality, leads to an uneven distribution of accessibility across population groups. A multi-objective optimization approach has been developed to identify optimal locations for bike-sharing stations by balancing demand coverage with spatial distribution (Nikiforiadis et al., 2021). This methodology takes into account urban infrastructure data, land use characteristics, and potential user demand to strategically position the stations while minimizing the need for redistribution. When applied to Thessaloniki, Greece, the study illustrates how various prioritization strategies, focusing either on demand or spatial coverage, affect the placement of the stations.

Geographic Information Systems (GIS) optimizes bike-sharing network expansion using spatial data to identify areas of high demand and prioritize resource allocation. Larsen et al. (2013) used a GIS-based approach to optimize the location of new cycling infrastructure using spatial data. The authors used a grid-cell technique to spatially aggregate data, such as demand indicators and existing infrastructure, which helps prioritize areas for development. In another study, Banerjee et al. (2020) identified the best locations for bike-sharing stations in Baltimore using a GIS-based approach that included over 1.6 million GPS data points from bike trips. Their method integrates location-allocation modeling with a modified version of Huff's gravity model. The GIS-based methodology effectively maps out optimal station placements that promise maximum accessibility and user engagement.

Imbalances in bike-sharing systems, such as underutilized stations or overburdened ones, can reduce efficiency and user satisfaction. To address this, Liu et al. (2015) proposed a method

that predicts station demand and maintains balance. Their study introduced a genetic algorithm-based optimization model designed to select high-demand locations. The model aims to maximize station usage and minimize imbalances by ensuring balanced pick-ups and drop-offs. A related study by Martinez et al. (2012) applied a heuristic algorithm to a Mixed Integer Linear Program. Their model optimizes bike station locations, fleet size, and the logistics of bicycle relocation activities within a day's operations. Their approach accommodates a mixed fleet of regular and electric bikes, integrating user demand, investment requirements, operational costs, and diverse fare structures.

Existing studies on optimizing the location of bike-sharing stations use discrete models, such as integer programming, location-allocation models, and heuristics, and primarily focus on accessibility and coverage, often neglecting the direct optimization of ridership potential. Such models may not fully capture network effects, which are necessary for modeling ridership. The following section will review Continuum Approximation algorithms to address facility location problems, specifically within urban transportation systems.

3.2.2 Continuum Approximation in Logistics

Discrete optimization models can provide optimal solutions for medium-scale logistics problems when sufficient computational power is available. However, since these problems are typically NP-hard, the computational effort required increases exponentially with size, making optimal solutions impractical for large-scale instances. The concept of Continuous Approximation (CA), initially proposed by Newell in the early 1970s (Newell, 1971, 1973), is an analytical alternative to discrete models for various logistics problems. CA models use continuous data representations and decision variables as density functions over time or space, simplifying the optimization process into manageable analytical operations. These models require less computational effort and help reveal managerial insights for large-scale practical problems (Ansari et al., 2018). This advantage has led to the application of CA models in problems related to facility location optimization, vehicle routing, and inventory management over the past years. Among the key developments in this area is the introduction of force-based CA models, which simulate interactions between facilities and demand fields to drive system design. In the facility

location problem proposed by Ouyang and Daganzo (2006), a CA-based algorithm was developed, which assumes spatially distributed customer demand and aims to identify a set of N terminals that minimize total logistics costs. This seminal study considers circular influence areas, each centered by a terminal to position it as closely as possible to its optimal location, drawing on Daganzo and Newell (1986), which found that such configurations yield near-optimum solutions. Ouyang and Daganzo (2006) used a force-based mechanism and adjusted terminal locations using impulsive forces to achieve equilibrium. The algorithm partitions a service area through a weighted Voronoi tessellation, ensuring each terminal is uniquely associated with one entire disk. The results demonstrate that the algorithm achieves feasible discrete terminal designs with costs closely aligning with theoretical predictions and validates CA cost estimates by comparing them to costs from actual designs.

CA has been used extensively in transit settings, examining diverse aspects and methodologies, including exploration of evenly distributed transit demand (Daganzo, 2010), as well as approaches that integrate grid and hub-and-spoke principles. Ouyang et al. (2014) considered a hierarchical system for space-dependent demand, further illustrating the versatility of CA in addressing complex transportation challenges. They presented a CA framework that varies bus route densities based on neighborhood demand: dense 'local' routes in high-demand areas and sparser 'main' routes in lower-demand areas. The design incorporates a power-of-two spacing scheme, enabling seamless integration of different route densities. The authors concluded that the iterative solution process of their CA model, supported by numerical experiments, offers enhanced adaptability and cost efficiency. This model allows transit networks to dynamically adjust to actual demand patterns—something traditional models, which assume uniform demand, struggle to do. As a result, their CA framework optimizes resource allocation and reduces operational costs, making it a valuable tool for urban transit planning.

A practical application of discretization methods for CA routing and distribution models has been demonstrated. Ouyang (2007) converts CA guidelines into discrete vehicle routing zones using a disk model and weighted centroidal Voronoi tessellations. This approach establishes initial zones with a disk model and refines it using weighted centroidal Voronoi tessellations to balance the delivery load across each zone. In another study, Caicedo et al. (2023) investigated

bike networks in cities with heterogeneous travel demands. The authors used CA techniques to optimize the layout of the bike-sharing network. Their approach balances user convenience and operational costs by determining bike lanes' optimal length and spacing, the number and location of bike stations, and the appropriate fleet size. They evaluated their model using data from the City of Guayaquil, Ecuador. The model's sensitivity analysis underscored the robustness of the proposed CA framework, demonstrating that, even under varying demand scenarios, the optimized network effectively supports cycling as an attractive and feasible mode of transport. A streamlined model was developed using CA techniques to determine the optimal strategic design variables for bike-sharing systems (Soriguera and Jiménez-Meroño, 2020). This approach focuses on optimizing components such as the size of the bicycle fleet, the number of stations, the total number of parking slots, and the required rebalancing intensity. Validated with data from the Bicing system in Barcelona, the model demonstrated economies of scale in bike-sharing systems.

Integrating the design of shared bike and public transit services, Luo et al. (2021) developed a CA model to optimize the system while considering it along a corridor with spatially heterogeneous demand. The model aimed to minimize the overall system cost by accounting for various transportation options available to users, including transit routes complemented by walking or biking and bike-only routes. The authors used a bi-level algorithm to solve the model: the upper level focuses on optimizing the hybrid network by leveraging analytical features of the CA model, while the lower level determines the route assignment equilibrium. The effectiveness of this CA model is validated through a case study involving an actual bus route in Chengdu, China.

In another study on shared mobility, the use of the CA model expanded from traditional static, single-facility location problems to dynamic EV sharing challenges that consider dynamic demands, origin-destination trips, and nonlinear vehicle charging times (Li et al., 2016b). This research aimed to optimize the locations of EV-sharing stations and appropriate fleet sizes to minimize total system costs, including station construction, vehicle charging, transportation, and fleet balancing under variable demands. The authors demonstrated how the CA methodology effectively manages large-scale heterogeneous challenges by breaking down the study

area into smaller neighborhoods. The findings indicated that CA provides an efficient and accurate estimate of total system costs, offering near-optimal solutions for the complex design of large-scale EV-sharing systems.

The existing literature on CA applications in transportation, particularly within shared mobility, primarily focuses on optimizing system costs or infrastructure placement, often overlooking the direct optimization of ridership potential in bike-sharing systems. Traditional models for bike-sharing location optimization mainly emphasize enhancing accessibility and often consider station coverage as the key factor in determining optimal station locations. However, these models may not sufficiently address the network effects to understand and enhance ridership dynamics, highlighting the need for more advanced algorithmic approaches. Furthermore, existing studies have not explored a force-based algorithm within a CA framework that incorporates both repulsion and attraction forces to maximize ridership, while allowing stations to continuously adjust their positions across the service area without being confined to a predefined network structure. Therefore, the primary objective of this study is to develop a CA model specifically designed to enhance ridership by optimizing the placement of bike-sharing stations. This approach departs from previous models by focusing on coverage and strategically adjusting each station's location within a service area where demand is spatially distributed according to origin-destination patterns, thereby maximizing ridership and network efficiency.

The next section introduces the force-based CA optimization framework in detail. We model bike-sharing stations as mobile agents influenced by demand-driven attraction and inter-station repulsion forces, iterating toward a spatial equilibrium that maximizes ridership across the network.

3.3 Problem Statement

Consider a set of N bike-sharing stations located at $X = \{x_1, x_2, \dots, x_N\}$ within a continuous service area S . Each station i is equipped with q_i docks and maintains an average of b_i bikes available, resulting in $q_i - b_i$ empty docks (where $q_i \geq b_i$ for all $i = 1, 2, \dots, N$). Here, b_i represents the station's capacity to supply bikes to users, while $(q_i - b_i)$ denotes its capacity to accept returned bikes. Both q_i and b_i are treated as parameters and remain constant throughout the

model.

The primary objective is to determine the optimal locations of the stations, X , to maximize the total ridership across the entire network. The origins and destinations are characterized by spatial density functions $\Lambda_I(x)$ and $\Lambda_J(x)$, respectively, for $x \in S$. We assume origins and destinations are independently distributed, a common assumption in CA models such as Daganzo (2010); Ouyang and Daganzo (2006), which is justifiable in highly populated urban settings. These functions define the spatial distribution of travel demand and serve as inputs to the attraction component of the force-based optimization, driving stations toward areas with higher ridership potential.

Each station i has an associated influence area A_i , defined as the region that encloses all origin and destination points closer to station i than to any other station. To ensure usability, the influence area of each station is bounded by a maximum allowable area, $A_{\max}$, such that $|A_i| \leq A_{\max}$ for all i . This constraint reflects the limited willingness of bike-sharing users to walk beyond a certain distance or time to access a station. Additionally, influence areas are designed to be mutually exclusive ($A_i \cap A_j = \emptyset$ for $i \neq j$), preventing overlapping regions that could lead to inefficient station utilization and imbalanced ridership distribution across the network.

The total trip production and termination from and to station i are defined as:

$$P_i(X) = \int_{x \in A_i} \Lambda_I(x),$$

and

$$T_i(X) = \int_{x \in A_i} \Lambda_J(x),$$

respectively. They capture the dependency of trip production and termination on stations set X , as station interactions and potential overlap between influence areas inherently link each station's to the others. Additionally, the demand densities represent aggregate trip generation and attraction intensities over a representative planning horizon. Consequently, the trip production and termination $P_i(X)$ and $T_i(X)$ reflect expected demand rates over this interval. To ensure the

operational reliability and efficiency of each bike-sharing station, it is imperative to prevent stations from being overwhelmed by demand that exceeds their handling capacities. Specifically, over the planning time interval, the total trip production P_i from station i must not surpass the number of available bikes, b_i , and the total trip termination T_i to station i must not exceed the number of empty docks, $q_i - b_i$. These constraints are defined as:

$$P_i(X) \leq b_i$$

$$T_i(X) \leq q_i - b_i$$

We define the ridership between the station pair (i, j) by $\lambda_{ij}(P_i, T_j, C_{ij}(x_i, x_j))$, which depends on the ridership production and termination, and the impedance factor, $C_{ij}(x_i, x_j)$ which depends on the station locations and their distance as a key feature of impedance. We explain the properties and empirical validation of the impedance function in the following sections.

We present the following mathematical problem to maximize the total ridership:

$$\text{Maximize}_X \quad \sum_{i=1}^N \sum_{j=1}^N \lambda_{ij}(P_i, T_j, C_{ij}(x_i, x_j)) \quad (3.1)$$

$$\text{s.t.} \quad x_i \in S, \quad \forall i \quad (3.2)$$

$$A_i \subseteq S, \quad \forall i \quad (3.3)$$

$$A_i \cap A_j = \emptyset, \quad \forall i \neq j \quad (3.4)$$

$$|A_i| \leq A_{\max}, \quad \forall i \quad (3.5)$$

$$\sum_{j=1}^N \lambda_{ij} = \sum_{j=1}^N \lambda_{ji} \quad \forall i \quad (3.6)$$

$$P_i(X) \leq b_i \quad \forall i, \quad (3.7)$$

$$T_i(X) \leq q_i - b_i \quad \forall i, \quad (3.8)$$

$$P_i(X) = \int_{x \in A_i} \Lambda_I(x) \quad \forall i, \quad (3.9)$$

$$T_i(X) = \int_{x \in A_i} \Lambda_J(x) \quad \forall i, \quad (3.10)$$

where Constraints (3.2) and (3.3) ensure the stations and their influence areas are enclosed within the service area, respectively, Constraint (3.4) ensures the influence areas are mutually exclusive, and Constraint (3.5) ensures the area of each influence area is less than the allowable limit. These constraints stipulates that the regions are not necessarily collectively exhaustive, especially when $A_{\max}N \ll |S|$. Constraint (3.6) ensures that for each station i , the total number of trips originating from station i is equal to the total number of trips terminating at station i . This flow equilibrium condition ensures a balanced fleet distribution across the network. This flow balance condition reflects a steady-state assumption commonly used in CA models for long-term network planning. While bike-sharing systems in practice often face temporal and spatial imbalances that require operational rebalancing, our model is intended for tactical-level design. By assuming average flow equilibrium over a representative planning horizon, we enable analytically tractable optimization of station placement and sizing. Operational interventions such as bike rebalancing or fleet management are considered outside the scope of this framework but remain important for short-term system performance. We further discuss this limitation in the conclusion.

Constraints (3.7) and (3.8) regulate the flow of incoming and outgoing trips for each station based on their respective capacities and fleet availability. Specifically, Constraint (3.7) guarantees sufficient bikes available to meet user demand for departures from each station, thereby preventing the overloading of stations and ensuring that stations can handle outbound demand. Similarly, Constraint (3.8) ensures that there are adequate empty docks to accommodate incoming bikes, thereby preventing situations where users are unable to return bikes due to a lack of available docks. These constraints collectively guarantee that each station i has sufficient

capacity to accommodate all incoming and outgoing trips without surpassing its operational limits.

3.4 Solution Methodology

The presented problem is a non-linear program with continuous and spatial (geometric) decision variables. The main continuous decision variables are λ_{ij} , and the spatial variables are $A_i(x), x \in S$. While the number of station location variables is modest, the problem is highly non-convex and tightly coupled: the objective function depends on inter-station distances, origin-destination demand fields, and dynamic station influence areas, all subject to non-overlapping constraints and capacity bounds. Such structural characteristics, especially the evolving geometry of influence areas and the ridership-maximizing objective make standard optimization techniques difficult to apply without relaxing the model's core components. This challenge is also noted in seminal CA studies where force-based heuristics were introduced to handle similar spatially continuous and geometrically constrained problems. Motivated by this, we use a force-based algorithm that iteratively adjusts station locations within a continuous space.

In this section, we first present a refined definition of the influence area and discuss how they relate to the constraints of the mathematical problem. We then present a new force-based algorithm that maximizes ridership by optimally moving the stations based on the applied forces until equilibrium is reached.

3.4.1 Influence Areas and Constraints

To implement the optimization model, we begin by formally defining the spatial influence areas of each station and how they relate to the key capacity and non-overlapping constraints. We define each influence area, A_i , of station i by a disk (similar to Ouyang and Daganzo (2006)) with center x_i and radius r_i , as follows:

$$A_i = \{x \in S \mid \|x - x_i\| < r_i\}. \quad (3.11)$$

This representation ensures that A_i is a subset of the spatial domain S where the influence of station i extends within a defined radius r_i . To prevent overlapping between influence areas based on Equations 3.4 we impose the following:

$$\|x_i - x_j\| \geq r_i + r_j, \quad \forall i \neq j. \quad (3.12)$$

This ensures that the distance between the centers of two disks is at least the sum of their radii, thus maintaining exclusive coverage for each station. Additionally, to adhere to the maximum allowable influence area, Equations 3.5, the radii r_i are bounded by:

$$r_i \leq r_{\max}, \quad \forall i. \quad (3.13)$$

This limitation restricts the radius to a feasible range, ensuring that each disk remains practical within the spatial domain S .

Leveraging the constraints on trip production and termination (Equations 3.7 and 3.8) along with the definitions of $P_i(X)$ and $T_i(X)$ (Equations 3.9 and 3.10), we derive bounds on the size of each influence area. Assuming the trip demand is approximately uniform within each influence area A_i , we locally approximate $\Lambda_I(x) \approx \Lambda_{I,\text{avg},i}$ and $\Lambda_J(x) \approx \Lambda_{J,\text{avg},i}$ for $x \in A_i$. With this station-level approximation, the constraints simplify to:

$$P_i(X) \approx \Lambda_{I,\text{avg},i} \cdot |A_i| \leq b_i, \quad (3.14)$$

$$T_i(X) \approx \Lambda_{J,\text{avg},i} \cdot |A_i| \leq q_i - b_i. \quad (3.15)$$

Substituting these approximations, the allowable area $|A_i|$ must satisfy:

$$|A_i| \leq \min \left(\frac{b_i}{\Lambda_{I,\text{avg},i}}, \frac{q_i - b_i}{\Lambda_{J,\text{avg},i}} \right). \quad (3.16)$$

The average density terms $\Lambda_{I,\text{avg}}$ and $\Lambda_{J,\text{avg}}$ are estimated based on initial area assignments.

For station i , these are defined as:

$$\Lambda_{I, \text{avg}, i} \approx \frac{P_i^{\text{initial}}}{|A_i^{\text{initial}}|}, \quad (3.17)$$

$$\Lambda_{J, \text{avg}, i} \approx \frac{T_i^{\text{initial}}}{|A_i^{\text{initial}}|}. \quad (3.18)$$

Here, P_i^{initial} and T_i^{initial} represent the initial trip production and termination values, respectively, and $|A_i^{\text{initial}}|$ is the initially assigned area for station i . Consequently, the size of A_i can be set as:

$$|A_i| \leq \min \left(\frac{b_i}{P_i^{\text{initial}}}, \frac{q_i - b_i}{T_i^{\text{initial}}} \right) \cdot |A_i^{\text{initial}}|. \quad (3.19)$$

Similarly, since we assumed the influence areas as disk and $|A_i| = \pi r_i^2$, we can write:

$$r_i = \min \left(\sqrt{\frac{b_i}{\pi P_i^{\text{initial}}}}, \sqrt{\frac{q_i - b_i}{\pi T_i^{\text{initial}}}} \right) \cdot r_i^{\text{initial}},$$

where $r_i^{\text{initial}} = \sqrt{|A_i^{\text{initial}}|/\pi}$. This formulation ensures that each influence area A_i is initialized within feasible bounds, balancing resource availability (b_i) and empty dock capacity ($q_i - b_i$) against trip demand densities over the planning time interval. By leveraging initial estimates, the model facilitates efficient iterative adjustments during optimization, ensuring that influence areas remain practical and operationally viable. With the influence areas defined, in the following we introduce a force-based algorithm that iteratively adjusts station locations in response to spatial demand and inter-station interactions.

3.4.2 A Force-based Algorithm

To optimize station locations, we use a force-based heuristic grounded in CA. The method models each station as a mobile facility that adjusts its position in response to spatial demand and its proximity to other stations. Without relying on a discrete candidate set, this approach enables stations to move freely within the continuous service area S to seek an equilibrium configuration. The algorithm initiates by defining each influence area, A_i , of station i by a disk

of area $A_{\max}$. This area is subject to three distinct repulsive or attraction forces influencing its optimal location. To prevent any part of a station's disk from crossing the service area boundary, we define a slightly contracted inner region $S' \subseteq S$, offset inward from the boundary of S . Stations are allowed to move freely within S' , ensuring that their entire influence areas remain contained within S . This assumption does not constrain the model's generality, as real-world cases with complex geometries can be addressed through various approaches, such as dividing them into connected components while preserving connectivity. The forces are detailed below:

Station-to-Station Force ($F_{SS_{ij}}$): The relationship between ridership and the distance between station pairs exhibits a log-normal distribution, as demonstrated by Toronto bike share data in Figure 3.1-a. This trend also appears in many other micro-mobility (bike-sharing specifically) services and has a logical justification in that when the trip distance is too short, the users may prefer to walk, and when it is too long, they may prefer a motorized mode. Hence, there is an optimal inter-station distance, D_{opt} , that maximizes ridership.

In our optimization model, the impedance function C_{ij} in objective formulation (Equation 3.1) represents the cost associated with traveling between stations. To align the force-based algorithm with the optimization model and to model the aforementioned relationship, we define the station-to-station force $F_{SS_{ij}}$ to reflect the impedance relationship captured by C_{ij} . $F_{SS_{ij}}$ exerts a repulsion force between stations that are closer than D_{opt} and an attraction force for those farther apart, encouraging station placements that approximate the optimal spacing and thereby maximize ridership, as illustrated in Figure 3.1-b. The station-to-station force is mathematically formulated as:

$$F_{SS_{ij}} = \log(P_i)\log(T_j) \left(\frac{1}{\exp\left(-\frac{(\ln(d_{ij}) - \ln(D_{opt}))^2}{2\sigma^2}\right)} - 1 \right) \cdot \text{sgn}(d_{ij} - D_{opt}), \quad (3.20)$$

where d_{ij} denotes the distance between two stations, P_i and T_j are respective production and termination of station i and j (Equations 3.9 and 3.10). σ is the standard deviation of the log-normal distribution derived from empirical data. The term $\text{sgn}(d_{ij} - D_{opt})$ determines the direction of the force vector, negative (repulsion) if $d_{ij} < D_{opt}$, and positive (attraction) if $d_{ij} > D_{opt}$.

This force directly relates to the impedance function C_{ij} in the optimization model, as it aims to position stations at distances that minimize impedance and thus maximize ridership between them. By moving stations toward D_{opt} , the algorithm effectively seeks to minimize the C_{ij} in the objective function. To prevent the excessive influence of stations with very high trip productions or terminations, we apply logarithmic transformations to P_i and T_j , which normalize their scales and ensure a balanced influence across the network. Additionally, to maintain computational stability and reflect practical considerations, we cap the station-to-station force beyond a specific maximum range derived from the data distribution. This capping aligns with limiting interactions between stations to a reasonable geographic scope, ensuring that excessively distant stations do not exert undue influence on each other. The functional characteristics of the station-to-station force are illustrated in Figure 3.1-c.

Accessibility Force $F_{AC_{ij}}$: In the optimization model, one of the key constraints is that stations must not overlap, ensuring that their influence areas A_i and A_j are disjoint. To enforce this non-overlapping constraint within the force-based algorithm, we introduce the accessibility force $F_{AC_{ij}}$. This force acts as a strong repulsion when the distance between two stations is less than a minimum allowable distance D_{min} , effectively preventing stations from clustering too closely together and violating the non-overlapping constraint. Mathematically, the accessibility force is defined as:

$$F_{AC_{ij}} = \begin{cases} K \cdot \left(\frac{D_{min} - d_{ij}}{d_{ij}} \right), & \text{if } d_{ij} < D_{min}, \\ 0, & \text{otherwise,} \end{cases} \quad (3.21)$$

where K is a constant determining the strength of the repulsion. Incorporating this force ensures that the stations maintain the required minimum separation distance, satisfying the non-overlapping constraint of the optimization model.




Figure 3.1: a) Ridership-distance(km) distribution, b) Station-station force concept, c) Station-to-station force function, d) Demand-to-station force function.

Figures 3.2-a and 3.2-b illustrate how the station-to-station force and accessibility force operate. Figure 3.2-a shows the influence areas (disks) of bike-sharing stations and how the accessibility force prevents overlapping by adjusting station positions through repulsion. Figure 3.2-b demonstrates the combined effects of attraction and repulsion from the station-to-station force acting on station number 1 from other stations, guiding it toward optimal spacing relative to D_{opt} .




Figure 3.2: Bike-sharing stations, disks, disk's radius, and service areas: a) Infeasible overlapping results in accessibility force, b) Station-to-station forces for station 1 in various scenarios, c) Demand-to-station forces for station 1.

Demand-to-Station Force (F_{O_i}/F_{D_i}): In our optimization model, maximizing total ridership involves strategically placing stations to capture high trip demand while satisfying capacity constraints. The trip generation and attraction potentials at specific demand points, P_O and T_D , apply attraction forces on the stations, denoted by F_{O_i} and F_{D_i} , respectively. These forces guide the stations toward higher-demand areas, thereby seeking to increase the trip productions and terminations defined in equations (1) and (2). The forces are formulated as:

$$F_{O_i} = \sum_{\text{all } O} \frac{P_O \cdot b_i^\beta}{d_{O_i}^\alpha} \quad (3.22)$$

and

$$F_{D_i} = \sum_{\text{all } D} \frac{T_D \cdot (q_i - b_i)^\beta}{d_{D_i}^\alpha} \quad (3.23)$$

where d_{O_i} and d_{D_i} are the distances from origin point O and destination point D to the station

i , respectively. P_O and T_D represent the trip generation potential at origin point O and trip attraction potential at destination point D , quantified as the expected rate of trips originating from or terminating at these points. The station's capacity to supply bikes, b_i , and its capacity to accept bikes, $q_i - b_i$, reflect its ability to accommodate demand and aid in satisfying the Equations 3.7 and 3.8. The exponents α and β are to be determined through empirical analysis, reflecting the effects of distance decay and station capacity on the attraction forces. These forces are modeled following the principles of the classical gravity model, which stipulates that the interaction between two entities decreases with increasing distance between them. In this context, demand points with higher trip generation or attraction potentials apply stronger attraction forces on stations, encouraging station placement closer to areas with higher demand. This function captures the natural tendency for stations to be drawn toward densely populated or highly frequented areas to maximize ridership. Furthermore, the inclusion of station capacity factors, b_i and $q_i - b_i$, in the force equations reflects that stations with greater capacity to supply or accept bikes are more strongly attracted to high-demand areas. This means that stations with more available bikes (higher b_i) are more strongly influenced by regions with high trip generation potential (higher P_O), and stations with more available docking space (higher $q_i - b_i$) are more strongly influenced by regions with high trip attraction potential (T_D). The functional dynamics of these attraction forces based on distance are displayed in Figure 3.1-d. The figure demonstrates how the magnitude of the force decreases as the distance (d_{Oi} or d_{Di}) increases, consistent with the gravity model's distance decay effect. Figure 3.2-c shows the attraction forces acting on station number 1 from different demand points. The figure highlights how stations are drawn toward areas with higher trip generation and attraction potentials, resulting in optimal placement that balances accessibility and capacity constraints.

Attraction forces pull stations toward areas with high concentrations of trip origins and destinations, while repulsion forces work to prevent overcrowding by maintaining a minimum distance between stations. The interaction between these opposing forces allows the stations to adaptively reposition themselves, ultimately reaching a balanced configuration that responds to demand while ensuring adequate network coverage. This equilibrium, driven by these forces, is essential for maximizing ridership across the system. In the next section, we will present the

algorithm that implements this mechanism and guides the system iteratively toward achieving equilibrium.

3.4.3 Algorithm Steps

Building on defined force components and spatial constraints, in this section we present a complete algorithm that implements the proposed model through an iterative process. Algorithm 1 consists of several stages designed to iteratively optimize the placement of bike-sharing stations. The objective is to maximize system ridership, guided by a doubly constrained gravity model.

In the initialization stage (Lines 1–6), the algorithm sets up the initial station locations and corresponding areas A_i . Each bike-sharing station's influence area is conceptualized as a circular disk, initially assigned a maximum radius, r_{max} or disk's maximum allowable size, A_{max} . This initial assignment ensures that all stations have the potential to cover the maximum allowable area without immediate constraints, facilitating an unbiased starting point for the optimization process. As the algorithm progresses, the influence radius of each station is recalibrated based on station capacity to supply and accept bikes, denoted by b_i (number of bikes) and $q_i - b_i$ (number of empty docks) as well as trip generation and termination potentials P_i and T_i . The recalibration of influence area ensures that it remains proportionate to the station's capacity and the local demand dynamics. For mathematical explanations see Section 3.4.1. This formulation inherently respect the capacity constraints in Equations 3.7 and 3.8 in the optimization model. By constraining influence area, the algorithm ensures that trip productions and terminations of the station remain within manageable limits, preventing overloading of stations and maintaining operational balance across the network. In this stage, the production P_i and termination T_i values for each station are computed using the given trip origin and destination density functions $\Lambda_I(x)$ and $\Lambda_J(x)$, respectively. These densities represent the spatial distribution of trip generation and attraction potentials within the service area. This initialization sets the stage for the iterative process.

In the next stage (Lines 7–23), the algorithm calculates the vector forces acting on each station, which have both magnitude and direction. First, for interaction forces between stations,

the distance d_{ij} between each station pair is computed. If the distance is below a minimum threshold $D_{\min}$, an accessibility force $F_{AC_{ij}}$ is applied to prevent clustering. If the distance is below a maximum threshold $D_{\max}$, the station-to-station force $F_{SS_{ij}}$ is computed. Otherwise, the station-to-station force is set to be zero. Second, the demand forces acting on each station are calculated. Each station's demand forces from all origin and destination points are combined and aggregated as vectors. This process accounts for the direction and magnitude of each force, enabling the algorithm to determine the final demand force that positions the station in areas of high demand.

In the third stage (Lines 24–34), the algorithm calculates the resultant force on each station, which directs the movement of the disks. This force is a combination of inter-station forces (including station-to-station and accessibility forces) and demand-to-station forces, balanced by a factor η . This parameter controls the balance between inter-station and demand-to-station forces. Higher values prioritize forces between stations thereby promoting a distribution that prioritizes station spacing optimization, while lower values emphasize demand attraction aligning station locations more closely with demand patterns. This parameter plays a critical role in shaping the station layout and reflects the policy trade-off between network coverage and access to high-demand areas. To adjust station positions, the algorithm uses a step size μ that ensures smooth convergence. At each iteration, a maximum allowable displacement is computed as an exponentially decaying function of the iteration count, reflecting the diminishing adjustment magnitude as the system approaches equilibrium. The actual step taken by each station is then scaled proportionally to this maximum, based on the magnitude of the station's resultant force relative to the strongest resultant force. This dynamic scaling approach enables controlled updates in early iterations and finer adjustments as convergence progresses, maintaining spatial integrity and promoting numerical stability. Convergence is defined as the point at which all station movements Δ_i fall below a small threshold, indicating that the system has reached a stable force-equilibrium. While we do not formally prove optimality, the equilibrium configuration exhibits no further incentive for movement and is interpreted as a near-optimal solution under the modeled dynamics. Once the step size is determined, station locations are updated based on the resultant force and μ , and the movement Δ_i for each station is recorded. Following

the position update, the production and termination values for each station are revised according to their new locations. The influence area A_i is then recalculated (Equation 3.19). To maintain consistency, A_i is constrained within predefined bounds ($A_{\min}$ and $A_{\max}$), and production and termination values are updated accordingly.

In the fourth stage (Line 35), trip distribution is performed using a doubly-constrained gravity model. Subsequently, the total ridership is calculated for the current station configuration, thereby evaluating the effectiveness of the station placements in achieving the objective of maximizing system-wide ridership while adhering to capacity constraints outlined in the optimization model. Finally, in the convergence check (Lines 36-37), the algorithm computes the average displacement Δ_{avg} of all stations. The iterative process continues, with station positions being continuously adjusted, until Δ_{avg} falls below a predefined threshold ϵ , indicating convergence. At this point, the system is considered to have reached a force-based spatial equilibrium, where stations are optimally aligned with demand patterns, and no further movement would significantly improve total ridership. Subsequently, the algorithm returns the optimal station locations $\{x_i^*\}$ that maximize system ridership (Line 38).

In the subsequent sections, first, we provide a computational example to evaluate the performance of the proposed model and algorithm in a controlled environment (Section 3.5). We then apply the algorithm to a real-world case study using the bike-sharing system in the city of Vancouver (Section 3.6).

The proposed model frames the placement of bike-sharing stations as a dynamic spatial optimization problem using a force-based CA framework. Unlike previous CA approaches focusing on coverage or cost minimization, our model aims explicitly to maximize ridership. It achieves this by allowing stations to iteratively adjust their positions in response to spatial demand gradients and interactions between stations. The adjustments incorporate both attraction and repulsion forces, taking into account origin-destination flows and total ridership as the overall system objective. This continuous-space formulation offers a scalable and analytically grounded alternative to discrete optimization, making it ideal for large-scale, data-driven network design. In the following section, we will demonstrate the model's applicability through computational examples with various demand distributions and conduct sensitivity analyses to

Algorithm 1 Station Placement Optimization Algorithm

Input: Initial station locations $X = \{x_1, x_2, \dots, x_N\}$, station docks $q_i = \{q_1, q_2, \dots, q_N\}$, available bikes $b_i = \{b_1, b_2, \dots, b_N\}$, density functions $\Lambda_I(x)$ and $\Lambda_J(x)$, step size μ with a step decay rate, force balancing factor η , convergence threshold ϵ , minimum and maximum influence areas A_{min}, A_{max} , minimum, maximum, and optimal station distances $D_{min}, D_{max}, D_{opt}$.

Output: Optimal station locations $\{x_i^*\}$ maximizing system ridership.

```
1:  $X \leftarrow \{x_1, x_2, \dots, x_N\}$ 
2: for  $i \in \{1, \dots, N\}$  do
3:    $A_i \leftarrow A_{max}$ 
4:    $P_i \leftarrow \int_{x \in A_i} \Lambda_I(x)$ 
5:    $T_i \leftarrow \int_{x \in A_i} \Lambda_J(x)$ 
6: end for
7: repeat
8:   for  $i \in \{1, \dots, N\}$  do
9:     for  $j \in \{1, \dots, N\}, j \neq i$  do
10:       $d_{ij} \leftarrow \|x_i - x_j\|$ 
11:      if  $d_{ij} < D_{min}$  then
12:         $F_{SSij} \leftarrow F_{ACij}$ 
13:      else if  $d_{ij} < D_{max}$  then
14:         $F_{SSij} \leftarrow \log(P_i) \log(T_j) \left( \frac{1}{\exp\left(-\frac{(\ln(d_{ij}) - \ln(D_{opt}))^2}{2\sigma^2}\right)} - 1 \right) \cdot \text{sgn}(d_{ij} - D_{opt})$ 
15:      else
16:         $F_{SSij} \leftarrow 0$ 
17:      end if
18:    end for
19:    for  $O \in \text{Origins and } D \in \text{Destinations}$  do
20:       $F_{O_i} \leftarrow F_{O_i} + \frac{P_O \cdot b_i^\beta}{d_{O_i}^\alpha}$ 
21:       $F_{D_i} \leftarrow F_{D_i} + \frac{T_D \cdot (q_i - b_i)^\beta}{d_{D_i}^\alpha}$ 
22:    end for
23:  end for
24:  for  $i \in \{1, \dots, N\}$  do
25:     $F_i \leftarrow \eta \cdot \sum_{j \neq i} F_{SSij} + (1 - \eta)(F_{O_i} + F_{D_i})$ 
26:     $x'_i \leftarrow x_i + \mu \cdot F_i$ 
27:     $\Delta_i \leftarrow \|x'_i - x_i\|$ 
28:     $x_i \leftarrow x'_i$ 
29:     $\Delta_{total} \leftarrow \Delta_{total} + \Delta_i$ 
30:     $P'_i \leftarrow \int_{x \in A_i} \Lambda_I(x), T'_i \leftarrow \int_{x \in A_i} \Lambda_J(x)$ 
31:     $|A'_i| \leftarrow \min(q_i / P'_i, (q_i - b_i) / T'_i)$ 
32:     $A_i \leftarrow \max(A_{min}, \min(A'_i, A_{max}))$ 
33:     $P_i \leftarrow \int_{x \in A_i} \Lambda_I(x), T_i \leftarrow \int_{x \in A_i} \Lambda_J(x)$ 
34:  end for
35:   $\lambda_{total} \leftarrow \sum_{i=1}^N \sum_{j=1}^N \lambda_{ij} \left( P_i, T_j, C_{ij}(x_i, x_j) \right)$ 
36:   $\Delta_{avg} \leftarrow \Delta_{total} / N$ 
37: until  $\Delta_{avg} < \epsilon$ 
38: return Optimal station locations  $\{x_i^*\}$ 
```

evaluate the model’s performance and robustness under different conditions.

3.5 Computational Example

We conduct a computational example involving a set of bike-sharing stations to evaluate the performance of the proposed algorithm. The study area is defined as a rectangle, representing a connected space with distinct boundaries. We add specific points as origins and destinations to represent demand distribution. The service area is partitioned into cells of dimensions dx and dy , and the distribution of origin and destination demands across these cells is modeled using density functions (similar to that explained in (Li et al., 2016b)). For instance, for origin demands, the density function was based on population density, reflecting areas from which bike-sharing users are likely to start their journeys. Conversely, destination demands were modeled using a density function based on employment and other work-related land uses, representing probable endpoints of these journeys. We initially set η at 0.5 to equally weigh the contributions from direct interactions between stations and the aggregate demands of origin and destination. Additionally, to ensure the algorithm produced realistic, stable, and interpretable outcomes across different demand patterns, we calibrated three key parameters: the demand attraction decay exponent (α), the station capacity responsiveness parameter (β), and the repulsion constant (K) governing inter-station spacing. These parameters were tuned through a structured sensitivity analysis across both the bi-focal and gradient demand scenarios. Specifically, we varied each parameter within a plausible range while observing model behavior with respect to convergence stability, final station configurations, and total ridership. The parameter α controls how quickly demand attraction decays with distance, and was set to reflect typical user willingness to walk or bike. The exponent β modulates the influence of station capacity on demand attraction, capturing diminishing returns for large stations. The constant K determines the strength of the repulsion force and was adjusted to enforce spatial separation without destabilizing the algorithm.

To validate the model, we consider various scenarios, including different demand distributions and their sensitivity analysis. We begin with a bi-focal demand pattern and three sensitivity analyses within this context. We conclude by examining a gradient demand pattern to assess

the model's robustness and adaptability across diverse demand distributions.

3.5.1 Bi-Focal Origin-Destination Demand Analysis

We characterize the bi-focal demand pattern by concentrating the demands around two central locations. This analysis demonstrates the model's performance in scenarios with highly concentrated demand distributions, reflecting urban environments where major amenities or employment hubs exert significant gravitational influences on transportation and mobility patterns. To construct this demand pattern, we synthetically place origin and destination points in two opposite quadrants of the service area, representing, for example, residential and employment zones. Each point is assigned a weight based on population or attraction intensity, and its influence is diffused across space using a Gaussian kernel to generate smooth spatial demand densities. The service area is segmented into a grid of cells, and each cell's origin and destination demand values are calculated based on proximity to these weighted points (based on demand density functions). This process enables controlled simulation of localized demand hotspots that reflect realistic urban dynamics. Figure 3.3-a depicts the density distribution of origin and destination demand. We focus on a network of 20 bike-sharing stations, which are randomly distributed between two demand zones. A maximum radius initially defines each station's influence area. Figure 3.3-b illustrates the distribution of origin and destination densities across the network, highlighted by the gradient in the background. A red dot marks each station, its corresponding ID is labeled, and the blue circle shows its initial influence area. Station-to-station and demand forces are calculated for each station. The resultant forces at each station can be determined by calculating these forces. These forces at each station and the resultant forces in the initial iteration are depicted in Figures 3.3-c, and Figure 3.3-d, respectively. In Figure 3.3-c, the vectors representing forces are not scaled to facilitate their visualization. Subsequently, the stations are relocated according to the computed resultant forces and step sizes. Their influence areas also adjust dynamically based on the new station positions. This process ensures that stations are better positioned to meet demand. Figure 3.3-e captures the movement of stations and the changes in their influence areas after one iteration within the service area.

Density Heat Maps




Figure 3.3: Configuration for an example of $N=20$, a)Origin and destination density heat-maps, b)Initial configuration, c)Station-to-station forces at iteration 1, d)Resultant forces at iteration 1, e)Station movements and updated influence areas at iteration 1.

Figure 3.4 displays the results of nine different iterations. As seen in the figure, the stations gradually shift toward positions that better reflect the underlying demand distribution, with their influence areas evolving simultaneously. Initially, the station positions and their areas of

influence are more dispersed, but as the iterations progress, they converge into more optimized locations. The influence areas adjust accordingly, reflecting changes in demand and proximity to other stations. The changes across these iterations highlight the dynamic nature of the algorithm, where station placement and coverage continuously improve with each step, leading to a more optimized configuration.




Figure 3.4: Station movement and influence area adjustments over nine iterations for an example of $N=20$. Each sub-figure (a) to (i) represents an iteration where stations are re-positioned and their influence areas are recalculated.

The iterative process is terminated once the movement of all stations falls below the threshold. After 274 iterations, the process ceased. Figure 3.5-a depicts the ridership analysis of the iterative process. The ridership trend shows a rapid increase in total ridership, reflecting the

adaptive repositioning of bike-sharing stations in response to demand patterns. As the iterations progress, ridership gains decrease, culminating in a plateau that suggests a state of equilibrium has been reached. The convergence behavior underscores the effectiveness of the relocation algorithm, which ceases once the average total movement of all stations remains consistently below the threshold. Furthermore, the final station positions, illustrated in Figure 3.5-b, demonstrate a distribution that aligns with high-demand zones, thereby maximizing accessibility and operational efficiency. This alignment supports increased ridership and enhances the overall sustainability of the bike-sharing system.

Figure 3.5-c presents an accessibility heat map showing each station's final optimized location in addition to an accessibility heat map. The accessibility heat map highlights regions within an assumed walkable distance from the stations, revealing areas with optimal coverage. The shades ranging from green to yellow quantify the ease of access to bike-sharing stations, with darker green regions signifying higher accessibility. High Accessibility Zones are areas close to the stations. These regions are well-served by the current station placements. Conversely, areas with lower accessibility scores (lighter yellow shades) suggest potential gaps in service. These areas could benefit from strategically adding new stations to enhance coverage (potential gaps). This analysis can be further enhanced by incorporating more contextual layers like bike road networks. Adding such layers can provide a better understanding of how the current infrastructure supports the bike-sharing system and help assess potential congestion points.




Figure 3.5: a)Total ridership trend, b)Optimized station locations, c)Accessibility heat map

Figure 3.6 depicts the application of the model to a smaller-scale scenario with 10 bike-sharing stations. This figure illustrates the dynamic process of station movement from Figure 3.6-a to e, each frame showing how each iteration progressively optimizes station placements in

response to demand patterns. Early frames exhibit an initial dispersal of stations, while subsequent frames show a gradual convergence towards more strategically optimized locations. Figure 3.6-f synthesizes and tracks these iterative adjustments, providing a composite view that compares the trajectories of stations. Different colors in this figure represent various stages in the iterative process, with lighter hues indicating final station locations, illustrating the model's effectiveness in stabilizing station distribution over time.




Figure 3.6: Station movement and influence area adjustments over nine iterations for an example of $N=10$, smaller-scale scenario. a) to e) Represents an iteration where stations are re-positioned and their influence areas are recalculated, f) Station movements tracking over iterations.

We perform sensitivity analysis to evaluate the performance and robustness of the model under various conditions. To facilitate a more detailed visualization and analysis, this sensitivity analysis is performed using the example that includes 10 bike-sharing stations. The sensitivity analysis includes exploring variations in optimal distance (D_{opt}), assessing the impact of balancing parameter (η) in integrating both station-to-station interaction forces and demand-to-station forces (guiding the disk movements of each station), and testing the model on larger scale problems.

Analysis of Optimal Spacing, D_{opt}

An integral part of our sensitivity analysis focuses on determining the optimal distance, D_{opt} , between bike-sharing stations in the model. To this end, we define the optimal distance directly as D_{opt} and adjust it using a scaling factor, θ . We conduct a sensitivity analysis by setting $D_{opt} = \theta \times D_{base}$, where D_{base} is the baseline distance historically considered optimal. We varied θ from 0.2 to 2.0 times D_{base} to examine the effects of both significant reductions and increases in station spacing while keeping all other model parameters constant to isolate the impact of changing D_{opt} . This range was chosen to explore significant decreases and increases in station spacing, reflecting a wide array of urban layouts, and aims to balance station spread to maximize accessibility while minimizing unnecessary overlaps and excessive walking distances.




Figure 3.7: Comparative analysis of station distribution and ridership dynamics across θ or D_{opt} values, a) to c) Spatial distribution of bike-sharing stations for varying D_{opt} values, d) A composite view of station distributions for all D_{opt} values, e) Trends in total ridership in relation to different D_{opt} settings.

Figure 3.7-a to c displays the final station locations under various θ settings, showing how different θ values influence station placement. These visuals facilitate a comparative analysis of the impact of varying optimal distances on station distribution. Figure 3.7-d further illustrates

these differences by presenting each configuration for every θ in various colors and a comparative manner, aiding in visualizing changes across the settings. To further quantify the influence of θ on the bike-sharing system's performance, we plot the total ridership as a function of various θ values to observe the relationship between optimal distance value and ridership levels. The resulting plot, as shown in Figure 3.7-e, demonstrates that as θ increases, ridership remains stable or increases slightly before dropping significantly. This trend suggests an optimal range for θ . Specifically, increasing θ enlarges the optimal distance D_{opt} relative to the demand distances and demand patterns. As a result, users may have to walk longer distances to access bikes, which can significantly reduce the system's convenience and usability. Additionally, such a spread can lead to gaps in coverage, ultimately decreasing the service level and impacting overall system effectiveness.

Balance between inter-station and demand-to-station forces, η

We evaluate the impact of η , which modulates the balance between inter-station and demand-to-station forces. By altering η , we effectively shift the focus of the resultant force: increasing η enhances the influence of inter-station forces, thereby promoting a distribution that prioritizes station spacing optimization. Conversely, decreasing η emphasizes demand forces, aligning station locations more closely with demand patterns and potentially improving accessibility. In scenarios where η is lowered to significantly enhance demand responsiveness, there is a tendency to concentrate stations in specific high-demand areas. While this may increase accessibility in these areas, it can reduce service equity across the broader network, as more remote or less densely populated areas might experience reduced service levels or a complete lack of access.




Figure 3.8: Comparative analysis of station distribution and ridership dynamics across η values, a) to c) Spatial distribution of bike-sharing stations for varying η Values, d) A composite view of station distributions for all η values, e) Trends in total ridership in relation to different η settings.

Figure 3.8 shows how changes in η impact the bike-sharing system. Displaying the optimized station locations under various η settings in Figures 3.8-a to c shows how different η values influence station placement. Figure 3.8-a depicts that station locations are heavily concentrated in high-demand areas when η is set to 0. This ensures that stations are easily accessible where user density is greatest. However, this can lead to station overcrowding in these hot spots. This strategy can inadvertently exacerbate equity issues within the transportation network. When η is increased to 0.25, Figure 3.8-b, the spread in station placement reflects a shift towards balancing demand density with broader area coverage. This slightly expanded station distribution helps mitigate any station's intensity, potentially enhancing service reliability. By extending the reach of the network to cover more geographic areas, this adjustment aids in addressing equity concerns within the bike-sharing system. At $\eta = 1$, Figure 3.8-c, station distribution is primarily based on spacing rather than direct demand. This can lead to under-service in high-demand areas. All configurations are consolidated in Figure 3.8-d, where distinct colors represent each

set to facilitate direct comparison.

Figure 3.8-e shows the interplay between the configuration of bike-sharing stations and the network's ridership. When station placement prioritizes high-demand areas exclusively ($\eta = 0$), ridership peaks. This maximization of ridership is consistent with the observed dense clustering of stations in these high-demand zones, as shown in Figure 3.8-a. As η increases and introduces a balance between demand concentration and optimal station spacing, ridership initially declines. This decrease reflects a transition from a purely demand-driven strategy towards a model considering spatial efficiency. Ridership then ascends again as η continues to increase, suggesting an optimal point where the balance between covering high-demand areas and maintaining station distribution maximizes user uptake. This recovery in ridership can be attributed to the network achieving a more strategic placement of stations that optimizes overall coverage without sacrificing access to high-demand zones. As η approaches 1, and the station distribution strategy becomes predominantly governed by geometrical considerations rather than demand responsiveness, ridership sharply declines. This underscores the drawbacks of a layout that prioritizes optimal spacing over demand-centric station placement.

Scalability

As the last phase of sensitivity analysis, we evaluate the scalability of our algorithm by applying it to different scales of bike-sharing station setups, ranging from 5 to 100 stations. This scaling test help to understand how the model performs under a broad spectrum of operational sizes and real world problems.

Figure 3.9 presents the relationship between the number of bike-sharing stations and the iterations needed to reach equilibrium. Despite the increase in iterations required—from 239 to 837 as the system expands from 5 to 100 stations—this trend confirms the model's robustness and capability to operate efficiently across large network scales. Computation time also increases with the number of stations, scaling from less than 3 minute for 5 stations to to well over forty minutes for 100 stations. The data illustrate that the model maintains consistent performance while convergence times increase, affirming its suitability for large urban bike-sharing networks.




Figure 3.9: Scalability Analysis Results; a)Relationship Between Number of bike-sharing Stations and Iterations to Equilibrium, b)Station locations distribution, N=50.

3.5.2 Gradient Demand Impact Analysis

In this section, we distribute demand more uniformly across the area, moving away from a bifocal approach. This method allows us to explore the impacts on station placement and overall ridership patterns.




Figure 3.10: Gradient demand scenario results, $D_{opt} = 200$, $\eta = 0.6$, a) Origin and destination demand heat maps, b) Initial configuration, b) Station movements and updated influence area in iteration 1, d) Final station locations, c) Total ridership trend.

Figure 3.10-a illustrates the origin and destination density heat maps across the region, clearly demonstrating how the demands have been distributed. The initial configuration of this scenario, including the distribution of origin and destination densities across the network, highlighted by the gradient in the background, is shown in Figure 3.10-b. Unlike focal scenarios, the gradient distribution lacks natural clustering of demand, which typically guides the initial placement and spacing of stations. Consequently, the absence of concentrated demand areas means that achieving an optimal equilibrium—where station locations perfectly balance demand across all points—requires more iterations. This increase in iterations is a reflection of the heightened complexity and the expanded search required to approach an optimal configuration. Figure 3.10-c captures the movement of stations and the changes in their influence areas after one iteration within the service area.

The optimized location of stations is visualized in Figure 3.10-d provides a clear overview of how stations are strategically redistributed to better match the gradient distribution of de-

mand. The depiction highlights the final station placements optimized for maximum coverage and minimal redundancy. The optimized station locations directly respond to the demand patterns, which is evident from the origin and destination data. The applied model successfully identifies and implements an equitable distribution of stations in the gradient distribution of demand, which typically does not exhibit natural clustering. This arrangement ensures that no area is oversaturated with resources or neglected, thus maximizing coverage and operational efficiency. These areas' adaptive size and placement demonstrate the model's capability to adjust parameters dynamically in response to sparse distribution patterns, underscoring the flexibility and responsiveness of the optimization process to varying urban dynamics.

Figure 3.10-e presents the ridership trend associated with the optimized station placements. This graph displays the changes in ridership over the iterations, capturing the increase in ridership as station locations are refined. The trend line indicates a gradual improvement in system utilization. This visual representation underscores the direct correlation between station optimization and ridership, providing quantitative evidence of the system's enhanced performance post-optimization.

Analysis of Gradient Demand Example

In this example, we conduct a sensitivity analysis on the parameter η to examine its influence on station placement and system performance within a bike-sharing system characterized by a gradient demand distribution. The aim is to discern the ideal settings that optimize both user accessibility and network efficiency.

Figure 3.11-a to c illustrates the distribution of bike-sharing stations for different values of η , reflecting how the balance between inter-station forces and demand-to-station forces alters station placement. When η is set to zero, reflecting a strong focus on demand intensity, all stations cluster at the single point of highest demand. This contrasts with the bifocal scenario, where the presence of two dominant demand peaks naturally attracts station clustering around these points, optimizing service for the most active parts of the network. As η increases, stations begin to disperse, extending coverage and supporting a more optimal distance between stations. As η approaches one, the placement strategy shifts too far towards optimal spacing without

sufficient regard to underlying demand patterns, potentially leading to under-service in areas that, while not at the peak, still maintain significant demand. All configurations are consolidated in Figure 3.11-d, where each setting is represented by distinct colours to facilitate direct comparison.

Figure 3.11-e displays the ridership trends corresponding to these configurations. When η is minimal, indicating a demand-focused distribution strategy, stations are concentrated at a single high-demand point, which results in lower overall ridership due to failure to support the optimal distance that maximizes ridership across the network. As η increases, there is an initial dip in ridership when stations begin to spread out, potentially extending beyond traditional high-demand zones but supporting a more optimal station-to-station distance. This spreading phase is important for finding a balance where system accessibility is maintained without compromising the convenience of having stations close to where users are located. However, as η approaches 1, the station distribution becomes predominantly driven by spacing rather than demand responsiveness, leading to a decrease in ridership. This decline underscores the importance of aligning station placement with actual user demand patterns and the optimal distance between stations to ensure the bike-sharing system's effectiveness. Moreover, comparing the sensitivity analyses of two different demand distribution patterns highlights that varying demand conditions may necessitate distinct η trend and values to best capture localized peaks while ensuring adequate coverage. In particular, this parameter inherently depends on both policy priorities and the underlying demand distribution, underscoring the need for context-specific calibration.




Figure 3.11: Comparative analysis of station distribution and ridership dynamics across η values in gradient demand scenario, a) to c) Spatial distribution of bike-sharing stations for varying η Values, d) A composite view of station distributions for all η values, e) Trends in total ridership in relation to different η settings.

3.6 Case Study

In this case study, we apply the model to the real-world scenario of the Vancouver bike-sharing system. This system comprises 245 stations distributed across the city of Vancouver, offering a substantial dataset for our analysis and model validation.

We source the station locations and capacity information from the General bike-sharing Feed Specification (GBFS) data. GBFS is an open data standard that provides a uniform method for sharing public bike-sharing system data. GBFS feeds typically include information on station status, bike availability, system alerts, pricing plans, or other details essential for comprehensive system management. We use this standard for the Vancouver bike-sharing system to obtain accurate data about system status and station specifics.

The model incorporates sociodemographic information and Points of Interest (POIs) to assess

demand density. We use census data, which provides comprehensive sociodemographic statistics, including population density, employment rates, and age demographics. Additionally, the Mapbox API delivers a rich dataset of POIs, enhancing our understanding of land use and specific points such as shopping centers and offices. The integration of these datasets enables us to develop a demand density function. We formulate this function by considering that population density predominantly influences origin demands, while employment hubs and recreational centers primarily drive destination demands. These factors are modeled to mirror the realistic distribution and intensity of both origin and destination demands within the urban landscape, allowing for precise calibration and validation of our bike-sharing model. Specifically, we construct spatial demand density functions by assigning population and POI-based demand indicators to Dissemination Areas (DAs). The DAs, represented by outlined polygons, serve as the smallest geographic units for which census data is available. These values are then spatially distributed within each DA to create continuous demand surfaces that reflect the local intensity of trip origins and destinations. These surfaces serve as model inputs for the force-based optimization.

Figure 3.12-a illustrates the spatial distribution of bike-sharing stations across various DAs in Vancouver. The bike-sharing stations are visualized as circles, with color and size reflecting their dock capacity. Stations with low dock capacity are depicted in green, medium dock capacity stations are shown in orange, and high dock capacity stations are represented in red. The radius of each circle scales with the station's dock capacity, providing a visual representation of how station density and dock capacity are distributed across the city.

Demand distribution across the DAs in Vancouver are shown in Figure 3.12-b to d. The origin demand reflects areas where trips are likely to begin, primarily based on population density and the distribution of housing units (Figure 3.12-b). The intensity of the colors represents areas with high densities of trip origins. In contrast, the destination demand focuses on likely trip endpoints, modeled using work-related POIs such as employment hubs, commercial centers, and other high-activity locations (Figure 3.12-c). This approach allows for a spatially-informed analysis of demand, emphasizing both the residential and commercial dynamics within Vancouver. Figure 3.12-d illustrate the spatial distribution of both origin and destination demands

across the study area.




Figure 3.12: a)City of Vancouver bike-sharing system, b)Origin demand density, c)Destination demand density heat map, d)Combined origin and destination demand heat map.

To determine the optimal distance D_{opt} between bike-sharing stations, we use ridership data provided by the Vancouver bike-sharing provider, Mobi. This dataset includes detailed trip information such as trip ID, start and end times, trip duration, bike type, user membership, and the origin and destination station names. Initially, missing information is corrected using other resources such as GBFS, and sparse missing values in ridership are removed to ensure data integrity. In addition to excluding outliers that indicate atypical trip durations, stations within the University of British Columbia campus area and their corresponding ridership data are omitted from the analysis. These stations are considered isolated, predominantly serving intra-campus travel, and thus do not reflect the broader dynamics of the city-wide bike-sharing system. Subsequently, we analyze trip frequencies relative to the distances between stations by

plotting these variables to generate a distribution graph. This analysis elucidates the relationship between ridership and station distance, revealing a peak where ridership is maximized. This peak corresponds to D_{opt} , the optimal distance between stations.

For the influence area of each station, we establish a maximum radius corresponding to a 400-meter walking distance and a minimum of 100-meter radius. These parameters allow for a realistic assessment of station coverage based on typical walking distances, ensuring stations are well-situated to meet local demand.

We explore the introduction of additional bike-sharing stations into the network considering three distinct expansion scenarios. Below outlines the findings from our implementation of scenarios, where we added 5, 10, and 20 new stations to the Vancouver bike-sharing network. Initially, the new stations were placed randomly within the city boundaries. Afterwards, the model aimed to identify their optimal locations, considering the existing stations' fixed positions. We applied the same calibration methodology described in the computational example to the Vancouver bike-sharing data. Specifically, we tested ranges of α , β , and K to ensure model stability, convergence, and realistic spatial layouts under real-world demand conditions. Through this process, we selected $\alpha = 1.2$, $\beta = 1.8$, and $K = 10^6$ as effective values that consistently yielded improved ridership and spatially balanced station distributions across the 5-, 10-, and 20-station expansion scenarios. While the optimal value of parameters may theoretically depend on factors like local station density or urban layout, we observed that the calibrated values remained effective across defined expansion scenarios. The configurations of the new stations for the five new stations scenario in iteration 1 is illustrated in Figure 3.13. The resultant forces and dynamic adjustments made by the algorithm during this iteration are depicted in Figure 3.13-a and b. Figure 3.13-c displays the final location of these new stations.



Figure 3.13: Five new station scenario, a)Resultant forces in iteration 1, b)Station movements in iteration 1, c)Final location of new stations.

Figures 3.14-a and 3.14-b illustrate the initial random locations of the stations along with their movements after the first iteration for the 10 and 20 station scenarios, respectively. Figures 3.14-c and 3.14-d, on the other hand, depict the optimized station locations for each scenario. These visualizations highlight the transition from the initial configurations to the optimized placements, demonstrating the effectiveness of the algorithm in different scenarios in improving station distribution. In both cases, the newly added stations are located toward the periphery of the existing network. This spatial outcome results from the interplay between repulsion forces from the fixed set of 245 stations, the tendency to maintain optimal inter-station distances, and

the attraction exerted by nearby demand concentrations. Together, these forces guide new stations toward areas that are under-served yet still demand-responsive, enabling balanced network expansion without violating spatial constraints.



Figure 3.14: a)Station movements-10 new station- Iteration 1, b)Station movements-20 new station- Iteration 1, c)Optimized location of 10 new stations, d)Optimized location of 20 new stations.

To further examine the implications of deployment strategy, we extend our analysis to include a full-network optimization scenario in which all stations are placed simultaneously from scratch, without treating any existing locations as fixed. This contrasts with the phased deployment cases previously analyzed, where new stations are added incrementally to an established network. By comparing these approaches, we evaluate the influence of early deployment decisions and the effects of path dependency on overall network performance and spatial coverage. This com-

parison is particularly relevant for long-term planning contexts, where infrastructure expansion typically unfolds in stages. To ensure comparability, we maintain consistent model parameters, including the value of η , across all scenarios. Figure 3.15 illustrates the resulting station layout under the full deployment approach. In this scenario the stations are more evenly distributed across the service area, reflecting an optimized response to the spatial demand patterns. This leads to improved accessibility and coverage, particularly in underserved areas. The degree of spatial dispersion and coverage equity can be modulated through policy-driven parameter tuning, most notably the force balancing factor η , which controls the trade-off between inter-station spacing and demand alignment. In practice, this allows planners to steer the resulting station layout toward specific objectives, such as maximizing overall ridership, promoting equitable service access, or enhancing network efficiency.




Figure 3.15: Optimized station configuration for the full deployment scenario, where all stations are placed simultaneously without any fixed infrastructure.

In terms of ridership performance, the full deployment scenario achieves a modestly higher total ridership compared to the phased approach. This improvement arises because, in the full deployment setting, all stations are free to reposition in response to demand patterns and station interactions, enabling a network-wide optimized configuration. In contrast, the phased

strategy optimizes only the newly added stations while treating previously deployed stations as fixed. As a result, the model needs to adapt to potentially suboptimal legacy placements, which can limit the achievable ridership. This trade-off underscores the practical tension between long-term system efficiency and the staged, path-dependent nature of real-world infrastructure deployment.

From a computational perspective, both deployment strategies rely on force-based updates driven by demand and interactions between stations, but their computational profiles differ significantly. In the phased strategy, existing stations remain fixed, which means that inter-station forces are derived only between the newly introduced stations and all other stations, both fixed and new. Similarly, demand-based forces are computed solely for the new stations, and only their positions are updated iteratively. This selective updating reduces the number of force evaluations needed and leads to faster convergence. In contrast, the full deployment scenario jointly optimizes the positions of all stations, necessitating the derivation of inter-station forces for every pair of stations, along with demand-based forces for the entire network. Consequently, the volume of force computations increases significantly, often requiring more iterations to reach equilibrium. Despite this, the full deployment remains computationally efficient. On a standard machine (Intel i7, 16GB RAM), the full optimization of all stations in Vancouver completes in approximately 91 minutes. In comparison, each phased addition completes in under 10 minutes. While the per-iteration cost remains manageable in both cases, the total runtime is generally longer for the full deployment scenario. Nonetheless, this approach is still significantly more scalable than discrete optimization methods, which must assess a combinatorially large set of potential site configurations and quickly become computationally infeasible in large-scale problems, such as in the Vancouver case. This chapter force-based algorithm guides stations toward optimal configurations through continuous spatial interactions.

3.7 Conclusion

This study presents a force-based continuum approximation algorithm designed to strategically place bike-sharing stations in order to maximize ridership. The algorithm adjusts station locations dynamically based on assessed forces, iteratively refining their positions to reach a near-

optimal configuration. It reliably converges to a stable spatial equilibrium that is robust and interpretable. In practice, this equilibrium reflects a solution in which stations are well-aligned with demand while maintaining appropriate spacing and respecting capacity constraints. Unlike traditional methods that primarily focus on coverage and often depend on discrete placement models—usually neglecting ridership and assuming a structured network—our approach prioritizes ridership within a continuous spatial framework. This allows stations to be relocated freely within the service area, ensuring they are positioned strategically to attract users and enable efficient trip completions. Computational experiments validate the efficacy of this algorithm. The results show that in the beginning, stations move significantly while searching for spots that meet the optimal distance and fulfill the attraction demands. As the iterations progress, the movement patterns of the stations reveal a complex interplay of factors that contribute to optimal station placement. These factors include the inter-connectivity between stations, their geographical centrality within the landscape, and their proximity to areas of high demand. This intricate balance ensures that each station is well-placed to serve current demand and positioned to enhance the overall cohesion and efficiency of the network.

The sensitivity analysis results in different demand distribution patterns highlight the adaptability of the algorithm in responding to varying demand densities and urban layouts. The results emphasize the importance of balancing station-to-station interactions and demand-to-station forces using parameter η which depends on both policy priorities and the spatial demand pattern. Such dual consideration prevents the excessive clustering of stations within overly congested areas, which can lead to saturation. It also avoids placing stations in low-demand areas where they would remain underutilized. By varying this balancing parameter and tracking changes in spatial layouts and total ridership, planners can align station placement with the specific demand profile and strategic goals. The sensitivity analysis further reveals that there is an optimal range for station spacing that, when achieved, maximizes ridership. Fine-tuning this spacing is essential for optimizing the performance. Additionally, assessing the algorithm's performance for larger problems demonstrates how the model performs under a broad spectrum of operational sizes and real-world challenges. In addition, we showcase the application of the model through a case study of the Vancouver bike-sharing system. The results highlight the model's

effectiveness in strategically expanding the network by identifying optimal locations for new stations. This approach ensures that each new addition enhances overall system efficiency and improves service, demonstrating the model's relevance to real-world urban mobility challenges. While this study primarily focuses on maximizing ridership, the integration of inter-station and demand-based forces indirectly consider other concerns such as equity by preventing station clustering and ensuring balanced spatial distribution across the service area. This approach helps to provide fair access to bike-sharing services while optimizing overall network performance. One modeling limitation in this study is the assumption of balanced inflow and outflow at each station, which reflects a long-term planning perspective rather than short-term operational fluctuations. Future research could incorporate dynamic demand patterns or rebalancing costs to address real-time imbalances. Another limitation is the assumption of fixed bike capacity and availability, which may not fully capture real-world variability. Addressing dynamic fleet behavior could enhance the model's applicability.

To the best of our knowledge, this study is the first to formulate a force-based CA model that explicitly targets network-wide ridership maximization in bike-sharing systems. By integrating both attraction forces, aligned with spatially distributed origin-destination flows, and repulsion forces, ensuring station separation and spatial balance, our model extends the CA literature beyond its traditional focus on cost or coverage objectives. This framework offers a tractable and adaptable approach to capturing demand-responsive station dynamics in a continuous-space setting. Future research should focus on determining the optimal number of bike-sharing stations, which could provide valuable insights for long-term strategic planning. Additionally, incorporating a diverse fleet that includes electric bikes could improve service efficiency and cater to diverse user needs. Exploring bike lane prioritization in station placement is another promising direction, as this can help ensure safer and more accessible routes for users, ultimately enhancing the overall effectiveness of the system.

Chapter 4

Electrification and Charging

Infrastructure

Abstract: The increasing popularity of electric-bike (e-bike) sharing systems has introduced new challenges in managing bike availability and optimizing charging infrastructure within station-based networks. The efficient placement of chargers improves service consistency, minimizes downtime, and maximizes ridership. This study proposes a unified framework that integrates State of Charge (SOC) dynamics, charger placement, and ridership optimization within a Markovian system for e-bike sharing networks. The system is modeled as a two-dimensional Markov chain, capturing the birth-and-death processes of low- and high-SOC bikes, along with charging transitions that influence system dynamics. We derive steady-state probabilities and use them to compute expected ridership, which serves as the objective in a station-level location optimization problem. In the proposed model, equipping a station with charging infrastructure electrifies all docks at that station, enabling immediate recharging of docked low-SOC bikes. We develop a heuristic algorithm coupled with the two-dimensional Markov chain to optimize charger placement and maximize ridership. To ensure scalability for large networks, we introduce a single-pooling state approximation that reduces computational complexity while preserving accuracy. We compare the heuristic's performance against exhaustive enumeration and find that it consistently achieves near-optimal solutions within 0.5-1.5% of the global optimum while reducing the computation time. Sensitivity analyses examine the effects of number of chargers, dock capacity, and charging rate on system performance. Comparative results across exact and approximate methods highlight the effectiveness of the proposed framework for scalable and accurate infrastructure planning in e-bike sharing systems. We validate the proposed framework using real-world case studies in Pittsburgh, Vancouver, and San Francisco, representing small,

mid-sized, and large bike-sharing systems. Across these cities, the model identifies high-impact charger locations based on SOC-aware ridership dynamics and demonstrates scalability under varying electrification scenarios. Sensitivity analyses show that as charger coverage increases, the algorithm shifts from targeting central high-demand hubs to broader network coverage, improving accessibility and reducing low-SOC occurrences. These results confirm the model's adaptability to different urban scales and mobility patterns.

Keywords: E-Bike Sharing Systems, Charger Placement Optimization, State of Charge (SOC) Dynamics, Ridership Maximization, Facility Location Problem.

4.1 Introduction

Sustainability in urban transportation has become a major focus for policymakers and practitioners worldwide, driving the growth of shared micro-mobility systems. Electric bikes (e-bikes) have become popular because they overcome hilly terrain, extend travel range through battery-assisted propulsion, and make cycling a more accessible and practical option for a wider range of users. For example, in New York City's Citi Bike program, e-bikes have increasingly been favored over conventional bikes, particularly for longer trips and in areas with challenging topography (Xie et al., 2024b). Overcoming these physical and distance-related limitations allows e-bikes to reduce dependence on private cars. Philips et al. (2022) estimate that replacing car trips with e-bikes could save up to 24.4 MtCO₂ annually, especially in car-dependent rural and suburban areas. In shared micro-mobility systems, e-bikes serve both as a complement to and a replacement for traditional bicycles. E-bikes have expanded cycling access by attracting users who might not have considered biking due to distance or physical exertion. Their growing popularity has significantly boosted the overall demand for shared cycling, with studies reporting a 65.9% increase in total ridership despite a 22.6% decline in traditional bike-sharing trips (Li et al., 2024). As e-bike sharing systems mature, they introduce new operational challenges, particularly around managing battery depletion and ensuring vehicle availability, necessitating infrastructure planning strategies tailored to their unique requirements.

Station-based bike-sharing systems are shifting toward e-bikes as users prefer electric over traditional bikes. In 2023, e-bikes accounted for 46% of all station-based bike-sharing trips in the U.S., despite making up only one-third of the total fleet (National Association of City Transportation Officials, 2024). This trend is particularly pronounced in major metropolitan areas. For instance, in Los Angeles, e-bikes were used eight times more frequently than traditional bikes, while in New York City, their usage was four times greater (National Association of City Transportation Officials, 2024). As e-bike usage grows, expanding both the fleet and charging infrastructure will be essential to meet rising demand. A well-devised expansion plan, along with reliable charging infrastructure, can significantly boost usage rates. For example, in Los Angeles, increasing the share of the e-bike fleet from 7% to 16% led to nearly a 300% rise in trip volume (National Association of City Transportation Officials, 2024). These trends emphasize the need for e-bikes to be accessible in bike-sharing systems, as their adoption greatly depends on fleet availability, efficient infrastructure, and streamlined operations to ensure reliable service and meet increasing demand.

E-bikes and other shared electric micromobility modes of travel face unique operational challenges, particularly regarding the need for regular recharging. As an e-bike's State of Charge (SOC) decreases during use, strategic planning is required to maintain operational availability. The charging process can be time-consuming and interrupt vehicle availability, directly impacting service reliability. Therefore, optimizing charging strategies is key to sustaining system performance and meeting user needs. Several charging options are available for e-bikes, including plug-in charging stations, battery-swapping systems, and hybrid solutions that combine multiple energy sources to enhance charging efficiency (Thomas et al., 2015). Each option has distinct advantages: plug-in stations are widely adaptable, battery-swapping systems provide faster turnaround times, and hybrid options improve energy management and operational flexibility.

Charger deployment at docking stations provides a more efficient and scalable way to manage battery levels in station-based e-bike sharing systems. Unlike free-floating e-bike fleets that often depend on manual battery swapping and redistribution, station-based systems can use fixed charging stations to streamline operations and enhance reliability. The deployment of charging

stations reduces labor costs of manual battery swapping or repositioning, eliminating the need for dedicated crews or independent contractors. Additionally, charging stations improve system efficiency by serving as centralized hubs where e-bikes with low-SOC can be replenished and redeployed as fully charged vehicles. However, establishing charging facilities in docked e-bike systems presents several challenges, including installation costs and the need to balance user demand with infrastructure availability. The location and number of charging stations are critical design factors; deploying charging stations at every docking point is often impractical due to the significant construction and maintenance costs involved. Therefore, the strategic planning of charging networks must consider these constraints while ensuring optimal performance and sustainability.

Markovian systems and queuing theory have been widely used in modeling shared mobility services to analyze station availability and system performance. Previous studies have examined bike availability management in station-based systems, focusing on inventory control and rebalancing strategies. One of the foundational works in this area is Raviv and Kolka (2013), which proposes an optimal inventory management for bike-sharing stations using Markovian queuing theory and optimizes station-level availability. In addition to operational inventory management models, empirical studies have shown that user accessibility and bike availability significantly impact ridership, emphasizing the importance of maintaining high availability levels to sustain system usage (Kabra et al., 2020). Some studies address the complex dynamics of shared networks, including fluctuating bike availability between stations and varying customer demand, by developing alternative modeling approaches. For instance, Li et al. (2016a) introduced a mean-field computational approach to approximate large-scale bike-sharing networks, while Datner et al. (2019) developed an inventory optimization framework accounting for station interactions and demand spillovers. However, these models do not account for SOC dynamics and assume that all bikes are interchangeable, an assumption that does not hold for e-bikes, as their availability is constrained by battery depletion and charging. Other studies have focused on fleet rebalancing (Fricker and Gast, 2016; Samet et al., 2018), exploring methods to redistribute bikes across the network; however, they do not explicitly consider charging restrictions, which significantly impact the usability and availability of e-bikes. To address charging limita-

tions, some studies have investigated battery swapping strategies (Yang et al., 2022; Zhao et al., 2024), which eliminate waiting time for recharging but introduce additional operational costs and logistical challenges. A growing body of research has focused on charging infrastructure optimization for electric micromobility. Some studies have proposed optimization frameworks for allocating shared e-bikes (Liu et al., 2017b), while others have developed models to optimize the placement of e-bike charging stations (Chen et al., 2020). Some studies have also explored the design of mixed bike-sharing systems, incorporating factors such as waiting time penalty (disutility) and social welfare considerations (Song et al., 2024). Various metaheuristic algorithms, including genetic algorithms and swarm optimization, are used to solve infrastructure planning and operational challenges. Liu and Ouyang (2025) developed an aspatial queuing network model to optimize charging station placement and fleet management for hybrid docked-dockless micromobility systems. Their framework integrates SOC transitions, user incentives for charging station drop-offs, and truck-based repositioning strategies, embedding these elements into a constrained nonlinear optimization model that minimizes system-wide costs. The exact non-linear programming solution in Liu and Ouyang (2025) relies on closed-form approximations of the queueing model, which makes it computationally efficient but less flexible in capturing state transitions of the granular system. Their model treats charging station utilization and vehicle repositioning as cost-driven decisions aimed at minimizing the total system cost.

Despite extensive research on bike-sharing operations, charging infrastructure planning, and stochastic modeling, there is currently no unified framework that integrates SOC dynamics, charger placement, and ridership optimization by capturing the stochastic dynamics of station states through a Markovian modeling approach. Existing studies primarily focus on cost minimization, fleet rebalancing, or infrastructure planning in isolation, without fully capturing the interdependent effects of charging events on system performance and user demand. This study aims to advance the state-of-the-art by proposing a unified framework that integrates SOC dynamics, charger placement, and ridership optimization within a Markovian system for e-bike sharing networks, explicitly accounting for charging events as part of the system’s evolution. To achieve this, we model station dynamics as a two-dimensional Markov chain, capturing the birth-death processes of low- and high-SOC bikes and charging transitions that influence

system performance. We reformulate the charging infrastructure planning problem as a station-based location optimization model, strategically selecting charger placements to maximize ridership while accounting for network-wide bike availability dynamics in an e-bike sharing system. To overcome computational challenges in large-scale networks, we introduce simplification methods that make it feasible to analyze complex steady-state dynamics while maintaining model accuracy. Our results demonstrate that these simplifications balance computational efficiency and accuracy, making the framework suitable for high-capacity implementations where exact solutions become computationally prohibitive. We first evaluate the proposed framework through computational experiments, demonstrating convergence properties, benchmarking against brute-force enumeration, and performing sensitivity analyses on key system parameters, including the number of chargers, station capacity, and charging rates. These experiments highlight how system performance responds to electrification levels and model simplifications. We then validate the framework’s applicability through real-world case studies in Pittsburgh, Vancouver, and San Francisco, representing small, mid-sized, and large bike-sharing systems. Across these cities, the model identifies high-impact charger locations under varying urban mobility patterns, station densities, and trip distance distributions. The findings confirm that the proposed framework effectively captures the operational complexities of e-bike sharing systems and offers a data-driven, scalable solution for sustainable micromobility infrastructure planning. The remainder of this chapter is organized as follows. Section 4.2 presents a comprehensive review of the relevant literature. Section 4.3 outlines the modeling framework, including the Markov Chain formulation and the associated optimization approach. Section 4.4 describes the solution methodology in detail and introduces simplification strategies to improve scalability for large-scale networks. In Section 4.5, we evaluate model performance and compare alternative methods through computational experiments. Section 4.6 presents a real-world case study to demonstrate practical applicability. Finally, Section 4.7 summarizes the key findings and discusses potential directions for future research.

4.2 Literature Review

The literature on shared mobility systems provides a foundation for understanding inventory management, fleet rebalancing, and charging infrastructure planning, but it often addresses these challenges in isolation, particularly for e-bike sharing systems where SOC dynamics play a critical role. This section reviews key studies in traditional bike-sharing systems and electric micromobility, highlighting their contributions and limitations in addressing ridership optimization, stations and charger placement, and the integrated challenges of SOC dynamics.

4.2.1 Inventory Management in Regular Bike-Sharing Systems

Inventory management in traditional bike-sharing systems is essential for ensuring bike and dock availability, minimizing user dissatisfaction due to empty or full stations. Markovian models and queuing theory are commonly used to capture station dynamics and optimize operations. Raviv and Kolka (2013) proposed a transient inventory modeling approach using a non-time-homogeneous Markov chain, modeling stations as double-ended queuing systems with finite capacity. They defined a user dissatisfaction function based on customer abandonment rates due to unavailability, proving its convexity to facilitate optimization, and introduced a discrete approximation method to set optimal inventory levels dynamically. Similarly, Fricker and Gast (2016) developed a stochastic model for homogeneous systems using M/M/1/K queues, focusing on minimizing "problematic stations" with no bikes or docks. Their results suggest an optimal fleet size of roughly half the station capacity, supplemented by user incentives (e.g., encouraging returns to less crowded stations) and redistribution mechanisms to balance bike availability. They also evaluated a single-pooling approach, aggregating stations into virtual nodes to simplify complexity. A demand-driven Markovian model is introduced by Liu et al. (2024) for bike rebalancing, capturing transient state transitions influenced by demand substitution effects across stations, enabling adaptive strategies for equitable bike distribution. In addition to queueing-based and Markovian modeling approaches, recent study by Freund et al. (2017) optimizes bike-sharing station inventory by reallocating dock capacity using multimodular function minimization, offering a scalable alternative to traditional queue-based models.

For large-scale systems, computational complexity poses a challenge due to fluctuating bike availability and varying demand. Li et al. (2016a) proposed a mean-field computational approach to simplify multi-dimensional Markov processes, approximating station behavior through the network's average dynamics using time-inhomogeneous queuing models and nonlinear birth-death processes. This method calculates steady-state probabilities for problematic stations, offering greater accuracy than simpler M/M/1/K models. Similarly, Samet et al. (2018) used a closed queuing network with a Repetitive Service Random Destination blocking mechanism, aggregating stations into virtual nodes and solving via Entropy Maximization. Tested on a 20-station system, this approach balanced accuracy and computational efficiency, providing reliable performance indicators like bike and dock availability.

Optimization and tactical planning aimed at improving system reliability and user experience have also been key research focuses. Datner et al. (2019) developed a framework for inventory optimization, considering station interactions where shortages create demand spillovers to neighboring stations. They used a guided local search algorithm to set inventory levels, minimizing excess travel time due to unavailability. Neumann-Saavedra et al. (2020) proposed a service network design framework, integrating resource management into tactical planning through optimized master tours for bike redistribution, using a time-expanded network and a tailored metaheuristic to minimize failed ride requests under budget constraints. Studies on broader shared vehicle systems provide additional context: Nair and Miller-Hooks (2011) developed a stochastic mixed-integer program with joint chance constraints for vehicle redistribution under demand uncertainty, He et al. (2021b) modeled customer preferences based on station connectivity and demand flows, and Banerjee et al. (2022) used a closed queuing network to optimize pricing and rebalancing across bike-sharing and other shared mobility platforms, focusing on throughput, revenue, and social welfare. While these studies advance inventory management, rebalancing, and optimization in traditional bike-sharing systems, they do not address the SOC dynamics and charging constraints critical to e-bike systems, highlighting the need for a more integrated approach.

4.2.2 Charging Infrastructure and Battery Management

Shared electric mobility services, such as e-bikes, face unique logistical and operational challenges due to battery management and charging needs, which reduce fleet availability during recharging. Effective management of SOC, strategic placement of charging infrastructure, and synchronization of charging with vehicle repositioning are critical for ensuring network efficiency. Recent research has used queuing network analyses, stochastic optimization, and combinatorial approaches to address these challenges, focusing on battery swapping, charging station deployment, and SOC-aware operations.

There have been studies focusing on optimizing battery swapping strategies and developing incentive mechanisms to encourage riders to drop off low-battery e-bikes at a subset of destination stations (Yang et al., 2022), or formulating a location-routing problem that optimizes the location of battery-swapping cabinets and the routing for technicians to visit depleted shared e-bikes (Zhao et al., 2024). Xie et al. (2024a) used nonstationary queuing models and fluid approximations to optimize battery inventory in a locker-based e-bike swapping system, reducing waiting times and abandonment rates, offering insights into dynamic operational management. While these strategies focus on dynamic operations, the deployment of charging stations in docked micromobility systems has emerged as a critical area of research, aiming to support scalable and efficient e-bike sharing networks. Chen et al. (2020) proposed a bi-level programming model to optimize the deployment of e-bike sharing stations, including determining the number of shared e-bikes and charging piles at each station. The model integrates an e-bike sharing network equilibrium in the lower level, which captures travelers' responses to station placement and charging demands, and a system planner's optimization objective in the upper level to minimize total system costs. The bi-level model is reformulated into a single-level mixed-integer program, solved using a hybrid particle swarm optimization algorithm. Wang et al. (2024) analyzed charging pile placement on a university campus using a spatiotemporal demand framework, leveraging trip chain models and accessibility metrics to prioritize high-demand locations while minimizing traffic disruption. In another study, a stochastic multiperiod optimization model is developed to determine the optimal deployment of traditional bikes and e-bikes within an integrated system (Zhu, 2021). This approach underscores the necessity of balancing system profitability with

user satisfaction and operational reliability.

SOC dynamics are central to e-bike system performance. Liu et al. (2017b) proposed a joint optimization framework for e-bike resource allocation, addressing station locations, dock assignments, and bike distribution to maximize service quality. Using Poisson models to predict user behavior, they introduced a greedy algorithm to iteratively allocate resources, reserving auxiliary resources for peak demand periods, with results from New York City's Citi Bike system showing improved service efficiency. Zhou et al. (2023) and Wang and Szeto (2021) used continuous-time Markov chains to model SOC dynamics, with the former optimizing battery swapping and rebalancing using a Markovian strategy to improve profitability, and the latter focusing on battery electric vehicle repositioning to minimize charging costs. In another study, Song et al. (2024) developed an optimization framework for mixed bike-sharing systems, focusing on the impact of SOC on availability. Their model integrates delay costs (e.g., when users cannot find charged e-bikes) to optimize station locations and capacities for maximum social welfare. The lower-level problem models user behavior via a multi-period, multi-modal network equilibrium, capturing interactions between travel modes, routes, and charging availability, while a Genetic Algorithm with decomposition techniques solves the problem, highlighting the need to balance user satisfaction, operational efficiency, and infrastructure costs. Managing resources in large-scale shared micromobility systems requires balancing vehicle availability, recharging constraints, and operational costs (Akturk et al., 2024). A Markov Decision Process framework has been developed to optimize vehicle repositioning and charging strategies while accounting for stochastic demand fluctuations and fleet dynamics. Akturk et al. (2024) also introduced a mean-field approximation technique to reduce computational complexity by replacing station-level randomness with deterministic system dynamics, ensuring scalability for large networks.

A hybrid system with both docked stations and dockless availability in shared micromobility networks introduce unique challenges in fleet management, charging infrastructure planning, and service reliability. In a recent study, Liu and Ouyang (2025) proposed an aspatial queuing network model designed to optimize charging station deployment and service operations for hybrid docked and dockless electric micromobility systems, specifically focusing on shared

e-scooters and e-bikes. They consider SOC dynamics, stochastic rider demand patterns, user pickup and drop-off behavior, charging station utilization, and truck-based repositioning operations. By embedding these dynamics within a constrained nonlinear optimization framework, the study jointly optimizes system parameters, including the optimal number, location, and capacity of charging stations, SOC-based user incentives, and truck-based battery repositioning strategies. The model emphasizes flow conservation across distinct vehicle states (idle, in use, charging) and uses steady-state queuing network analysis to evaluate system performance. To address the optimization complexity, the authors implemented a decomposition approach that combines enumeration of discrete decision variables, gradient-based optimization techniques for continuous variables, and iterative refinement based on queuing network equations. Numerical experiments presented in the study demonstrate substantial improvements in operational efficiency, cost reductions, and enhanced rider accessibility achieved through the strategic integration of charging stations in mixed docked-dockless micromobility systems.

Studies on related systems, such as e-scooters and electric vehicles, provide additional context. Chen et al. (2018) optimized charging and battery-exchange station locations for e-scooters, using a dual-objective model to balance costs and service coverage. Their framework distinguishes between Battery Charging Stations and Battery Exchange Stations, integrating factors such as land prices, service distance constraints, and station capacities. The study uses a Multi-Objective Particle Swarm Optimization algorithm to identify optimal station placements and allocations. Lai and Li (2024), Barth and Todd (1999), Brandstätter et al. (2017), and Roni et al. (2019) addressed charging infrastructure and fleet operations for electric ride-hailing and car-sharing systems, focusing on downtime, scalability, and cost-efficiency. Integrated frameworks combining charging and repositioning include He et al. (2021a) and Chen and Liu (2023), which optimized charging infrastructure and vehicle operations for EV sharing systems using queuing networks and stochastic programming, highlighting the importance of balancing long-term planning with real-time decisions.

While the literature has made significant strides in addressing inventory management, fleet re-balancing, battery management, and charging infrastructure planning for shared mobility systems, a critical gap remains in integrating these elements into a cohesive framework tailored to

the unique dynamics of e-bike sharing networks. Existing studies often treat SOC dynamics, charger placement, and ridership optimization as separate challenges, focusing predominantly on cost minimization or operational efficiency in isolation. For instance, Markovian models have been extensively applied to traditional bike-sharing systems, yet they rarely account for the battery constraints inherent to e-bikes ridership. Similarly, charging infrastructure studies provide valuable optimization frameworks but overlook the interdependent effects of charging events on system-wide performance. This fragmented approach limits the ability to fully capture the operational complexities of e-bike sharing systems, where battery management directly influences availability and ridership. This study advances the field by proposing a unified framework that integrates SOC dynamics, charger placement, and ridership optimization within a Markovian system and optimization model for e-bike sharing networks. Unlike prior work, we explicitly model charging events as part of the system’s evolution. We use a two-dimensional Markov chain to capture the interplay between low- and high-SOC bikes and charging transitions. We reformulate charger placement as a station-level location optimization model that maximizes system-wide ridership under capacity and SOC constraints. To ensure scalability, we introduce computationally efficient simplification methods that balance model accuracy and computational tractability, making the approach practical for large-scale implementations. This contribution bridges the gap between operational and infrastructure planning and provides a data-driven, scalable solution to enhance service reliability and support sustainable urban mobility.

4.3 Methodology

Consider an e-bike sharing network consisting of a set of stations denoted by $\mathcal{S} = \{1, 2, \dots, S\}$. Each station $s \in \mathcal{S}$ has a docking capacity C_s , which is the maximum number of bikes it can accommodate. A subset of the stations $\mathcal{S}_n \subseteq \mathcal{S}$ with the cardinality of N are equipped with chargers. We assume that all its docks support charging when a station has a charger, which is consistent with the operations of realistic services like Toronto and Vancouver. The assumption implies that any docked e-bike in a charging station can be charged. We seek to determine the optimal selection of charging stations to maximize overall ridership. The challenge lies

in strategically locating chargers across the network to reach the desired objective, taking into account the dynamics of e-bike movements in the network and the rate of charging.

The SOC of an e-bike represents its battery level, which directly influences its availability within a shared mobility system. Throughout its operation, an e-bike's SOC changes based on usage patterns, with battery depletion occurring during trips and recharging when docked at a charger-equipped station. In a general setting, e-bikes could be classified into multiple SOC levels at each station, ranging from fully depleted to fully charged, with intermediate states representing varying levels of battery range. To maintain computational tractability while still capturing the essential operational characteristics of battery depletion and charging dynamics, we adopt a two-state SOC representation of low-SOC and high-SOC. This assumption does not affect the generality of the model to more than two SOC levels. Low-SOC bikes have a depleted battery level but can still be rented, albeit with a lower preference for selection due to their limited range. If docked at a charging-equipped station, they become eligible for recharging. In contrast, high-SOC bikes are fully or sufficiently charged, making them more likely to be rented due to their greater range. These bikes do not require charging when docked.

When an e-bike completes a trip and is returned to a station, its SOC can vary depending on factors such as trip duration, user behavior, and battery consumption patterns. Some bikes may retain enough charge to remain in the high-SOC category, while others may drop to low-SOC, making them candidates for recharging if a charger is available on drop-off.

4.3.1 Markov Chain Model

We model each station $s \in \mathcal{S}$ as a Markov chain, with its state defined by the tuple (i_s, j_s) , where i_s and j_s are the number of low-SOC and high-SOC bikes at station s , respectively. The total number of bikes at a station must satisfy the following capacity constraint:

$$i_s + j_s \leq C_s, \quad 0 \leq i_s, j_s \leq C_s. \quad (4.1)$$

The Markovian system evolves over time, where state changes occur due to three types of events: arrivals, departures, and charging. All events occur according to independent Poisson

processes with specified rates (in units of events per time). Hence, the time between successive events of each type is exponentially distributed. The events transitioning between states are defined as follows:

Arrival Events (Inflow or Returns): Bikes returning to station s increase the number of either low or high-SOC bikes. A returning bike is classified as low or high-SOC based on its state at the end of the trip. The inflow rates are denoted by λ_s^L and λ_s^H , representing the arrival rates of low- and high-SOC bikes at station s , respectively. The transitions are feasible only if $i_s + j_s \leq C_s$ and can be written formally as

$$\begin{cases} (i_s, j_s) \rightarrow (i_s + 1, j_s) & \text{at rate } \lambda_s^L \quad \text{if } i_s + j_s < C_s, \\ (i_s, j_s) \rightarrow (i_s, j_s + 1) & \text{at rate } \lambda_s^H \quad \text{if } i_s + j_s < C_s. \end{cases} \quad (4.2)$$

Departure Events (Outflow or Rentals): Users rent bikes from station s , decreasing the count of available bikes. The rental process is demand-driven, meaning that a bike is taken only if it is available at the station. The rental rates are $\mu_s^L \times i_s$, and $\mu_s^H \times j_s$, where μ_s^L and μ_s^H denote the rental rates per bike for low- and high-SOC states, respectively. Thus,

$$\begin{cases} (i_s, j_s) \rightarrow (i_s - 1, j_s) & \text{at rate } \mu_s^L i_s \quad \text{if } i_s \geq 1, \\ (i_s, j_s) \rightarrow (i_s, j_s - 1) & \text{at rate } \mu_s^H j_s \quad \text{if } j_s \geq 1. \end{cases} \quad (4.3)$$

Charging Events: At charging-equipped stations ($s \in \mathcal{S}_n$), low-SOC bikes are charged, transitioning to the high-SOC state at a rate $\gamma_s \times i_s$, where γ_s represents the charging rate per low-SOC bike at station s . Specifically, γ_s is expressed as the probability that a low-SOC bike become a high-SOC state during its stay at a charging-equipped station. For charger-free stations $\gamma_s = 0$, meaning that bikes arriving with low-SOC remain in that state until they are rented again. Hence,

$$(i_s, j_s) \rightarrow (i_s - 1, j_s + 1) \quad \text{at rate } \gamma_s \times i_s \quad \text{if } i_s \geq 1, \gamma_s \neq 0. \quad (4.4)$$

Each of the aforementioned events results in a transition from one state (i_s, j_s) to another. The probability of transitioning from state (i_s, j_s) to a neighboring state is proportional to the rate




Figure 4.1: State transition diagram for a station with capacity $C_s = 3$. State (i, j) represents i low- and j high-SOC bikes.

at which the corresponding event occurs. These transitions form a structured two-dimensional Markov process, where changes in one state variable depend on interactions with the other. Figure 4.1 illustrates the transition diagram for a station with a capacity of $C_s = 3$, depicting all possible state transitions and providing a visual representation of how the system evolves. Each node (i_s, j_s) represents a state where the station holds i_s low-SOC and j_s high-SOC bikes. Since $i + j$ can not exceed 3, the valid states form a triangular region in the (i_s, j_s) plane, omitting any states for which $i_s + j_s > 3$. As a result, states along the boundary, such as $(0,3)$ or $(1,2)$, do not permit further arrivals in that respective dimension (i.e., low-SOC or high-SOC), since the station is already at capacity. Conversely, states with $i_s = 0$ (or $j_s = 0$) do not permit departures of low-SOC (or high-SOC) bikes. The directed edges indicate transitions caused by different events: blue arrows represent low-SOC bike arrivals at rate λ_s^L , green arrows correspond to high-SOC bike arrivals at rate λ_s^H , red arrows denote bike departures, with rates proportional to the number of available corresponding bikes (low- or high-SOC bikes), and dashed orange arrows indicate charging transitions, where low-SOC bikes become high-SOC at rate γ_s .

We define the state space at station s as $\mathcal{K}_s = \{(i_s, j_s) \mid 0 \leq i, j \leq C_s, i + j \leq C_s\}$ and K is the total number of states. The total number of feasible states for a single station is the number of

ways to choose nonnegative (i_s, j_s) such that $i_s + j_s \leq C_s$. For each $i_s = 0, 1, \dots, C_s$, the variable j can take values from 0 up to $C_s - i_s$. Thus, the total number of feasible states is:

$$K = \sum_{i_s=0}^{C_s} ((C_s - i_s) + 1) = 1 + 2 + \dots + (C_s + 1) = \frac{(C_s + 1)(C_s + 2)}{2}. \quad (4.5)$$

The total number of feasible states, given by $K = \frac{(C_s+1)(C_s+2)}{2}$, grows quadratically with the capacity-yielding 6 for $C_s = 2$, 10 for $C_s = 3$, and 15 for $C_s = 4$ - thereby exponentially increasing the complexity of the system as the state space expands with larger capacities.

Steady-State Formulation: To analyze the long-run behavior of the system, we formulate the steady-state equations for the two-dimensional Markov process for a station. Let $\pi_s(i_s, j_s)$ denote the steady-state probability of observing i_s low-SOC bikes and j_s high-SOC bikes at a given station, where $0 \leq i_s + j_s \leq C_s$. These probabilities satisfy two key conditions: (i) the total inflow into each state (i_s, j_s) equals the total outflow from (i_s, j_s) , and (ii) all state probabilities sum to one. For each feasible state (i_s, j_s) , the global balance equation is given by:

$$\begin{aligned} & \underbrace{\lambda_s^L \pi_s(i_s - 1, j_s)}_{\text{inflow from } (i_s-1, j_s), \text{if } (i_s \geq 1)} + \underbrace{\lambda_s^H \pi_s(i_s, j_s - 1)}_{\text{inflow from } (i_s, j_s-1), \text{if } (j_s \geq 1)} + \underbrace{\mu_s^L(i_s + 1) \pi_s(i_s + 1, j_s)}_{\text{inflow from } (i_s+1, j_s), \text{if } (i_s+1 \leq C_s - j_s)} \\ & + \underbrace{\mu_s^H(j_s + 1) \pi_s(i_s, j_s + 1)}_{\text{inflow from } (i_s, j_s+1), \text{if } (j_s+1 \leq C_s - i_s)} + \underbrace{\gamma_s(i_s + 1) \pi_s(i_s + 1, j_s - 1)}_{\text{inflow from } (i_s+1, j_s-1), \text{if } (i_s+1 \leq C_s - j_s+1, j_s \geq 1, \gamma_s \neq 0)} \\ & = \pi_s(i_s, j_s) \underbrace{\left[\underbrace{\lambda_s^L + \lambda_s^H}_{\text{if } (i_s + j_s < C_s)} + \underbrace{\mu_s^L i_s}_{\text{if } (i_s \geq 1)} + \underbrace{\mu_s^H j_s}_{\text{if } (j_s \geq 1)} + \underbrace{\gamma_s i_s}_{i_s \geq 1, \gamma_s \neq 0} \right]}_{\text{total outflow}}, \end{aligned} \quad (4.6)$$

Each term on the left-hand side corresponds to a transition arriving at state (i_s, j_s) from one of its neighboring states, while the right-hand side represents the total transition departing state (i_s, j_s) . The parameters λ_s^L , λ_s^H , μ_s^L , and μ_s^H denote the arrival and rental rates for low and high-SOC bikes, respectively, while γ denotes the charging rate for low-SOC bikes at a charging station. By convention, any infeasible state (i_s, j_s) where $i_s + j_s > C_s$ or either i_s or j_s is negative is assigned as $\pi_s(i_s, j_s) = 0$, ensuring that probability mass is distributed only among valid states. For notational simplicity, we will write $\pi_s(i, j)$ instead of $\pi_s(i_s, j_s)$ throughout the rest

of the paper, whenever the station index s is clear from context.

The complexity of solving steady-state probabilities grows with the station capacity C_s , as the state space size (Equation 4.5) increases quadratically, rendering analytical solutions intractable for large C_s . While closed-form solutions may exist for small C_s , no general pattern based on the C_s emerges as C_s grows, due to the system's high dimensionality and interdependent balance equations. Charging transitions complicate the Markov chain's structure, with probabilities depending on bike availability and charging dynamics (e.g., $\gamma_s i_s$), precluding product-form solutions. Appendix B provides the transition-probability matrix for small capacities, highlighting persistent complexity. Thus, rather than seeking explicit formulas for $\pi_s(i, j)$, we solve balance equations numerically using methods like power iteration and Gaussian elimination, which efficiently handle the system's scale and structure.

4.3.2 Optimization Model for Charger Placement

The charger placement problem in an e-bike sharing system is formulated as an optimization model to maximize ridership while accounting for station capacities, rental and return dynamics, and charging constraints. Each station $s \in S$ with maximum dock capacity of C_s is modeled as a finite-state Markov process with a two-dimensional state space, where the number of low-SOC and high-SOC bikes changes due to arrivals, departures, and charging events. The key decision involves selecting a subset of stations to equip with chargers, capturing the complex interdependencies in bike movement and state transitions throughout the network.

We introduce a binary decision variable $x_s \in \{0, 1\}$ to indicate whether station s is equipped with a charger. If $x_s = 1$, all docks at station s become charging-enabled, allowing low-SOC bikes to recharge to high-SOC at a rate γ . We measure ridership via the expected outflow of bikes across all stations. The expected outflow at station s is split into low-SOC and high-SOC components:

$$R_s^L = \sum_{(i,j) \in \mathcal{K}_s} \pi_s(i, j) \mu_s^L i_s,$$

$$R_s^H = \sum_{(i,j) \in \mathcal{K}} \pi_s(i,j) \mu_s^H j_s,$$

where μ_s^L and μ_s^H denote the respective rental rates. To accurately model the movement of bikes between stations while considering variations in their SOC, we assume that the expected outflows at each station are determined by the SOC transition matrix and the origin-destination (O-D) flow distribution. Specifically, we define:

$$\theta_{s \rightarrow t}^{XY} = Q_{s \rightarrow t} P_s^{XY}, \quad \forall s, t \in S, \forall XY \in \{LL, LH, HL, HH\}, \quad (4.7)$$

where $Q_{s \rightarrow t}$ represents the O-D fraction, indicating the proportion of bikes rented from station s that are ultimately returned to station t . The term P_s^{XY} is derived from the SOC transition matrix as below:

$$P_s^{XY} = \begin{pmatrix} p_s^{LL} & p_s^{LH} \\ p_s^{HL} & p_s^{HH} \end{pmatrix}, \quad (4.8)$$

where each entry of P_s^{XY} represent the probability a bike transitioning from SOC state X to state Y while traveling. For instance, p_s^{LL} denotes the probability that a low-SOC bike remains in a low-SOC state during transit, while p_s^{HH} represents the probability of a high-SOC bike retaining its charge. In many practical implementations, on-road charging is unavailable, resulting in $p_s^{LH} = 0$. The parameter $\theta_{s \rightarrow t}^{XY}$ thus quantifies the proportion of bikes moving from station s to station t while undergoing specific SOC transitions. This parameter is essential for accurately estimating the incoming bikes at each station (λ_t^L and λ_t^H), ensuring a realistic representation of bike redistribution dynamics across the network and define as below:

$$\begin{aligned} \lambda_t^L &= \sum_{s \in S} \left(R_s^L \theta_{s \rightarrow t}^{LL} + R_s^H \theta_{s \rightarrow t}^{HL} \right), \quad \forall t \in S \\ \lambda_t^H &= \sum_{s \in S} \left(R_s^L \theta_{s \rightarrow t}^{LH} + R_s^H \theta_{s \rightarrow t}^{HH} \right), \quad \forall t \in S \end{aligned}$$

The objective is to maximize the total ridership across all stations and we formulate the mathe-

matical optimization problem as below:

$$\text{Maximize } \sum_{s \in S} (R_s^L + R_s^H) = \sum_{s \in S} \sum_{(i,j) \in \mathcal{K}_s} \pi_s(i,j) (\mu_s^L i_s + \mu_s^H j_s) \quad (4.9)$$

$$\text{s.t } \sum_{(i,j) \in \mathcal{K}_s} \pi_s(i,j) = 1, \quad \forall s \in S \quad (4.10)$$

$$\gamma_s \leq x_s \gamma, \quad \forall s \in S \quad (4.11)$$

$$\sum_{s \in S} x_s = N \quad (4.12)$$

$$R_s^L = \sum_{(i,j) \in \mathcal{K}_s} \pi_s(i,j) (\mu_s^L i_s), \quad \forall s \in S \quad (4.13)$$

$$R_s^H = \sum_{(i,j) \in \mathcal{K}_s} \pi_s(i,j) (\mu_s^H j_s), \quad \forall s \in S \quad (4.14)$$

$$\lambda_t^L = \sum_{s \in S} (R_s^L \theta_{s \rightarrow t}^{LL} + R_s^H \theta_{s \rightarrow t}^{HL}), \quad \forall t \in S \quad (4.15)$$

$$\lambda_t^H = \sum_{s \in S} (R_s^L \theta_{s \rightarrow t}^{LH} + R_s^H \theta_{s \rightarrow t}^{HH}), \quad \forall t \in S \quad (4.16)$$

$$\sum_{(k,\ell) \in \mathcal{K}_s} \pi_s(k,\ell) T_s((k,\ell) \rightarrow (i,j)) = \pi_s(i,j) \sum_{(r,u) \in \mathcal{K}_s} T_s((i,j) \rightarrow (r,u)), \quad \forall (i,j) \in \mathcal{K}_s, \forall s \in S \quad (4.17)$$

$$T_s((i, j) \rightarrow (r, u)) = \begin{cases} \lambda_s^L, & \text{if } (r, u) = (i+1, j) \text{ and } i+j < C_s \\ \lambda_s^H, & \text{if } (r, u) = (i, j+1) \text{ and } i+j < C_s \\ \mu_s^L i, & \text{if } (r, u) = (i-1, j) \text{ and } i \geq 1 \\ \mu_s^H j, & \text{if } (r, u) = (i, j-1) \text{ and } j \geq 1 \\ \gamma_s i, & \text{if } (r, u) = (i-1, j+1) \text{ and } i \geq 1, (i-1) + (j+1) \leq C_s \\ 0, & \text{otherwise} \end{cases} \quad \forall s \in S \quad (4.18)$$

$$\pi_s(i, j) \geq 0, \quad \lambda_s^L, \lambda_s^H \geq 0, \quad x_s \in \{0, 1\}, \quad \forall (i, j) \in \mathcal{K}_s, \forall s \in S \quad (4.19)$$

The optimization objective aims to maximize the total expected ridership across all stations. The ridership is captured by the expected outflows of low-SOC and high-SOC bikes, reflecting the overall system utilization and demand satisfaction. To ensure that the Markov model maintains valid probability distributions, Constraint (4.10) enforces that the sum of steady-state probabilities equals one at each station. This guarantees a well-defined state representation.

The integration of chargers into the system directly influences the charging dynamics. When a station is equipped with a charger, the charging rate at that station is set to γ , as formulated in Constraint (4.11). This dependency ensures that only selected stations contribute to battery replenishment. Constraint (4.12) restricts the number of stations that can be equipped with chargers to exactly N .

The rental dynamics of bikes at each station are governed by Constraint (4.13) and (4.14), which quantify the expected outflow of low-SOC and high-SOC bikes, respectively. These equations incorporate steady-state probabilities and rental rates to estimate the number of departures. Maintaining a balance between inflows and outflows is essential for system stability. Constraint (4.15) and (4.16) ensure that each station receives an appropriate number of incoming bikes, calculated based on outgoing flows from other stations and transition probabilities of each pair of stations (Equation 4.7). These equations reflect the interdependencies between

stations and enforce network-wide equilibrium.

The fundamental principle and global balance equations of steady-state equilibrium is captured by Constraint (4.17), which ensures that the probability distribution at each station remains consistent. $T_s((k, l) \rightarrow (i, j))$ mathematically define the state transitions which represent how bikes move between different SOC states at a station and ultimately determine the station's steady-state distribution, as discussed in Section 4.3.1. These transition rates directly influence the steady-state probabilities $\pi_s(i, j)$ at each station by dictating how frequently bikes enter and leave different states. In particular, the balance equation (Constraint (4.17)) ensures that for each state, the total inflow rate equals the total outflow rate, maintaining equilibrium over time. Constraint (4.18), defines the transition rate function as discussed in Section 4.3.1 and define how bikes move between low-SOC and high-SOC states due to rentals, returns, and charging. Finally, the feasibility of the optimization problem is ensured through non-negativity constraints in Constraint (4.19), which enforce non-negativity on probabilities and bike flow variables while also restricting the decision variables x_s to binary values.

This optimization problem poses significant computational challenges due to several key factors. First, the Markovian state representation at each station introduces a large state space, as the number of possible system states grows exponentially with capacity. Second, the model incorporates nonlinear interdependencies between stations, particularly through the SOC transition dynamics and the coupling of inflows and outflows across the network. These dependencies create complex feedback loops where the state of charge at one station influences availability and demand at others, complicating the optimization landscape. Additionally, network-wide coupling in bike availability means that decisions regarding charger placements and station capacities impact the entire system, further increasing the problem's dimensionality and computational burden. Due to these challenges, solving the problem exactly using traditional optimization solvers may be computationally infeasible for large-scale systems, motivating the use of heuristic approaches, which will be explored in subsequent sections.

4.4 Solution Methodology

The optimization model in Section 4.3.2 establishes an interdependent system in which the steady-state probabilities $\pi_s(i, j)$ at each station depend on the overall network inflows $\{\lambda_s^L, \lambda_s^H\}$, which themselves depend on the rental outflows from every other station. Given this complexity and size of the optimization problem, exact methods may be computationally infeasible. Therefore, an iterative algorithm including ridership derivation based on equilibrium computation and ridership optimization is proposed to efficiently explore the solution space and identify near-optimal charging station configurations.

4.4.1 Ridership Derivation: Equilibrium Computation for a Fixed Charger Configuration

Suppose we are given a fixed subset of stations $S_n \subseteq S$ to be equipped with chargers. Let $\gamma_s = \gamma_{\text{charging}}$ if $s \in S_n$ and 0 otherwise. We aim to find the steady-state solution—that is, the equilibrium inflows $\{\lambda_s^L, \lambda_s^H\}$ and steady-state probabilities $\{\pi_s(i, j)\}$ for all stations s .

We begin by initializing the arrival rates at each station with a guess, denoted by $\lambda_s^{L,(0)}$ and $\lambda_s^{H,(0)}$. In each subsequent iteration $m = 1, 2, \dots$, we first calculate the steady-state probabilities for every station, $\pi_s^m(i, j)$, using the current inflow parameters $\lambda_s^{L,(m-1)}$ and $\lambda_s^{H,(m-1)}$. We obtain these probabilities numerically using methods such as power iteration and Gaussian elimination. Once the probabilities are determined, the expected outflows of low-SOC and high-SOC bikes are calculated for each station, $R_s^{L,m}$ and $R_s^{H,m}$ (Equations 4.13 and 4.14).

These outflows are then distributed across the network according to the origin-destination fractions and SOC transition probabilities (see (4.7)), yielding new inflow rates at each station t as $\hat{\lambda}_t^{L,m}$ and $\hat{\lambda}_t^{H,m}$. To improve numerical stability and guarantee convergence under typical monotonicity conditions, we adopt the Method of Successive Averages. Specifically, at iteration m , we update the inflow rates using a diminishing step size $\beta_m = 1/m$ so that

$$\begin{cases} \lambda_t^{L,m} = \beta_m \hat{\lambda}_t^{L,m} + (1 - \beta_m) \lambda_t^{L,(m-1)}, \\ \lambda_t^{H,m} = \beta_m \hat{\lambda}_t^{H,m} + (1 - \beta_m) \lambda_t^{H,(m-1)}. \end{cases} \quad (4.20)$$

This iterative averaging approach helps mitigate oscillations and is a standard technique in transportation and network flow problems, ensuring the algorithm progressively “smooths” out inflow estimates from earlier iterations. The iterative process continues until the change in the inflow rates, as measured by $\|\lambda^m - \lambda^{(m-1)}\|$, falls below a prescribed tolerance ϵ . At convergence, the total network ridership is computed by

$$\mathcal{R}^m = \sum_{s \in \mathcal{S}} (R_s^{L,m} + R_s^{H,m}),$$

and the set of steady-state probabilities $\{\pi_s(i, j)\}$, $\lambda_s^{L,m}$, and $\lambda_s^{H,m}$ are recorded as the equilibrium solution corresponding to the current charger configuration.

Consistency and Convergence. For each station s , define the local feasible set

$$\mathcal{F}_s = \left\{ (\pi_s, \lambda_s^L, \lambda_s^H) \mid \pi_s \in \Delta_{|\mathcal{K}_s|}, 0 \leq \lambda_s^L, \lambda_s^H \leq \Lambda_{\max} \right\},$$

where $\Delta_{|\mathcal{K}_s|}$ indicates that the state probabilities $\{\pi_s(i, j)\}$ are nonnegative and sum to one. Let $\mathcal{F} = \prod_{s \in \mathcal{S}} \mathcal{F}_s$ be the overall feasible region for the inflow rates and state probabilities at all stations.

When inflows $\{\lambda_s^{L,(m-1)}, \lambda_s^{H,(m-1)}\}$ are given, we solve the linear balance equations to obtain the steady-state probabilities $\{\pi_s^m\}$. Conversely, once $\{\pi_s^m\}$ are determined, we update inflows $\{\lambda_s^{L,m}, \lambda_s^{H,m}\}$ using

$$(\pi_s^{(m+1)}, \lambda_s^{L,(m+1)}, \lambda_s^{H,(m+1)}) = f(\pi_s^m, \lambda_s^{L,m}, \lambda_s^{H,m}), \quad \forall s \in \mathcal{S},$$

where $f : \mathcal{F} \rightarrow \mathcal{F}$ implements the balance equations and the inflow/outflow consistency relationships. Because these updates are all linear and affine transformations that preserve nonnegativity and normalization, the iteration remains in $\mathcal{F}$. Subsequently, this process converges to a

stable point where each station's inflows, outflows, and steady-state probabilities are consistent with one another. We refer to this stable point as the system's equilibrium for the fixed charger configuration.

4.4.2 Simplification Methods for Ridership Derivation in Large-Scale Problems

Single-Pooling State Approximation

For large station capacities C_s , solving the full two-dimensional Markov chain governing the steady-state distribution of low and high-SOC bikes, $\pi_s(i, j)$, in ridership derivation part of the algorithm becomes computationally intractable due to the state space scaling as $O(C_s^2)$. To address this challenge, we use the single-pooling state approximation, which simplifies the problem by tracking only the total number of bikes in the station, $k = i + j$, rather than separately tracking low-SOC (i) and high-SOC (j) bikes. This reduction transforms the original two-dimensional process into a more tractable one-dimensional Markov chain.

In this approximation, we assume that, on average, a fraction α of the k bikes are in low-SOC, while the remaining fraction, $(1 - \alpha)$, are in high-SOC. The key assumption is that α can be estimated based on system parameters and iteratively refined. By collapsing the state space from (i, j) pairs to a single count k , the system behaves as an $M/M/1/C$ queue. The arrival rate of bikes into the station is $\lambda = \lambda_s^L + \lambda_s^H$, the sum of low- and high-SOC bike returns. The service rate (rental rate) depends on the number of available bikes. Since a fraction α of bikes are low-SOC and $(1 - \alpha)$ are high-SOC, the effective rental rate when there are k bikes in the station is given by:

$$\mu(k) = k[\alpha\mu_s^L + (1 - \alpha)\mu_s^H] = kv. \quad (4.21)$$

This process can be visualized as a birth-death chain (Figure 4.2-a), where birth transitions ($k \rightarrow k + 1$) occur at rate λ when $k < C_s$, and death transitions ($k \rightarrow k - 1$) occur at rate vk when $k > 0$, where $v = \alpha\mu_s^L + (1 - \alpha)\mu_s^H$. The utilization factor, ρ , represents the ratio of the total bike arrival rate to the effective rental rate and is defined as:




Figure 4.2: Birth-death process representations for simplified 1D Markov chain, a)Single-pooling state approximation, b)Instant charging approximation

$$\rho = \frac{\lambda}{v} = \frac{\lambda_s^L + \lambda_s^H}{\alpha \mu_s^L + (1 - \alpha) \mu_s^H}. \quad (4.22)$$

A higher ρ indicates a busier system with more frequent station occupancy, while a lower ρ suggests a lower probability of full station states. The steady-state probability $\pi_s(k)$ for the total number of bikes follows the standard $M/M/1/C$ solution:

$$\pi_s(k) = \begin{cases} \frac{\rho^k}{\sum_{m=0}^{C_s} \rho^m}, & \text{for } \rho \neq 1, \\ \frac{1}{C_s + 1}, & \text{for } \rho = 1. \end{cases} \quad (4.23)$$

Considering the flow balance equations, the fraction of low-SOC bikes is determined by equating inflows and outflows. The inflow consists of returning depleted bikes at rate λ_s^L . The outflow occurs due to rentals at rate $\mu_s^L i$ and charging transitions at rate γi . At steady state, the expected number of low-SOC bikes remains constant, leading to the approximation:

$$\lambda_s^L \cong (\mu_s^L + \gamma)(\alpha E[k]). \quad (4.24)$$

This formulation assumes an averaged steady-state approximation rather than an exact conservation equation for all system states. Since charging transitions introduce variability in SOC levels, α is not deterministic but an expectation. Solving for α :

$$\alpha = \frac{\lambda_s^L}{(\mu_s^L + \gamma)E[k]}, \quad (4.25)$$

where $E[k]$ is the expected number of bikes as $E[k] = \sum_{k=0}^{C_s} k \pi_s(k)$. To obtain a consistent estimate of α , we use an iterative refinement approach. We begin with an initial estimate $\alpha^{(0)}$,

then compute the effective service rate

$$v(n) = \alpha^{(n)} \mu_s^L + (1 - \alpha^{(n)}) \mu_s^H$$

and determine the updated utilization factor as $\rho^{(n)} = \lambda / v(n)$. Using the updated ρ , we compute the steady-state distribution $\pi_s(k)$ from Equation (4.23) and evaluate the expected number of bikes. The estimate of α is then updated as

$$\alpha^{(n+1)} = \frac{\lambda_s^L}{(\mu_s^L + \gamma)E[k]}.$$

This process is repeated until convergence, ensuring that $|\alpha^{(n+1)} - \alpha^{(n)}| < \epsilon$ for a small threshold ϵ . Once we obtain the steady-state distribution $\pi_s(k)$, we approximate the distribution of low- and high-SOC bikes using a binomial model:

$$\pi_s(i, j) = \pi_s(k) \binom{k}{i} \alpha^i (1 - \alpha)^j, \quad \text{where } i + j = k. \quad (4.26)$$

This approach is motivated by the assumption that each bike independently has a probability α of being in low-SOC. The binomial model ensures that: 1) the expected fraction of low-SOC bikes matches α , 2) the total probability mass is correctly distributed across possible (i, j) states, 3) variability in SOC composition is realistically incorporated rather than assuming a fixed ratio. The single-pooling state approximation significantly reduces computational complexity, from $O(C^2)$ states to $O(C)$, making it feasible to analyze stations with large capacities. The iterative estimation of α ensures adaptability to different system parameters. Moreover, the comparison between this approximation and exact probabilities highlights its strong performance in capturing system dynamics while providing computational efficiency (Section 4.5.7). While some approximation errors exist due to the assumptions of constant α and independent SOC states, the method remains a robust and scalable approach for steady-state analysis in large-scale e-bike sharing systems.

Upper bound via Instant Charging Approximation

To establish a theoretical upper bound on station performance, we consider the limiting case where $\gamma \rightarrow \infty$, meaning that charging occurs instantaneously upon bike return. Under this assumption, any bike arriving to a charging station in a low-SOC state (i) is immediately charged to a high-SOC state (j), eliminating low-SOC bikes in steady state.

This assumption reduces the system to a single state of high-SOC bikes, transforming the original two-dimensional Markov chain into a one-dimensional birth-death process governed by j , the number of high-SOC bikes at the station. The station thus follows an $M/M/1/C$ queue with arrival rate of $\lambda = \lambda_s^L + \lambda_s^H$, as all returned bikes are immediately available in high-SOC. The service rate is $\mu_s^H j$, since all rentals occur at the high-SOC rate. This birth-death chain is illustrated in Figure 4.2-b. The steady-state probability distribution for the number of high-SOC bikes j follows the standard $M/M/1/C$ form and similar to single-pooling state approximation where the utilization factor is defined as $\rho = \frac{\lambda}{\mu_s^H}$. Since no low-SOC bikes exist in this limit, the probability distribution simplifies to:

$$\pi_s(i, j) = \begin{cases} \pi_s(j), & i = 0, \quad j \leq C_s, \\ 0, & i > 0. \end{cases} \quad (4.27)$$

This idealized scenario provides an upper bound on station performance by eliminating charging delays. While unrealistic, it serves as a benchmark for evaluating the potential impact of increasing charging rates. Furthermore, this limiting case underscores the role of charging infrastructure in ensuring bike availability. When γ is sufficiently large, the system effectively operates as if all stations are fully stocked with high-SOC bikes, removing state-dependent rental constraints. This insight is particularly relevant for optimizing charger deployment and capacity planning.

Although this approach provides a method to reduce the two-dimensional system to a one-dimensional representation, it imposes significantly stronger assumptions than the single-pool state approximation. Nevertheless, comparing real-world performance against this upper bound facilitates model validation and informs infrastructure improvements.

4.4.3 Ridership Optimization: A Heuristic Search for Charger Configuration

With the equilibrium solver established, we now seek an effective way to select which N stations should host chargers in order to maximize overall ridership. Instead of enumerating all $\binom{|S|}{N}$ possible configurations (which is computationally infeasible for large $|S|$), we implement a local-improvement or “swap” heuristic. We assess and potentially improve the initial charger configuration by identifying candidate stations for swapping. This involves evaluating the performance of charger-equipped stations to detect any that are underperforming, and simultaneously identifying non-charger stations that might achieve higher ridership gains if provided with a charger.

We consider three distinct approaches and corresponding trigger metrics to guide the selection of potential swaps in the charger configuration. Each station is assigned a score that quantifies its potential contribution to network ridership if its charging status were modified. These approaches are as follows:

Direct Ridership Improvement (DRI): This approach evaluates the marginal change in total ridership when the charging status of a station s is altered. The score is computed as:

$$\Delta \mathcal{R}_s = \frac{\partial}{\partial \gamma_s} \sum_{s \in S} (R_s^L + R_s^H) \approx \mathcal{R}(\gamma_s = \gamma_{\text{charging}}) - \mathcal{R}(\gamma_s = 0), \quad (4.28)$$

where γ_s indicates whether station s is equipped with a charger. A higher positive $\Delta \mathcal{R}_s$ suggests that adding a charger at station s could significantly increase ridership, while a negative value indicates that reallocating the charger elsewhere may be more beneficial.

Low-SOC Prevalence: Here, the focus is on quantifying how frequently a station encounters low-SOC bikes. For each station s , the score is defined based on the fraction of time the station holds at least one low-SOC bike and the ratio of rentals. The score is computed as:

$$\text{Score}_s^{(A)} = \sum_{\substack{(i,j) \in \mathcal{K}_s \\ i \geq 1}} \pi_s(i, j) \cdot \frac{R_s^L}{R_s^L + R_s^H}. \quad (4.29)$$

This formulation captures both the likelihood of low-SOC occurrences and their impact on

ridership. A high score indicates that station s frequently experiences low-SOC conditions, suggesting that installing a charger could mitigate unmet demand.

Charging Potential: In this approach, we estimate the additional number of high-SOC bikes that station s could supply if it were forced to have a charger. The score is given by:

$$\text{Score}_s^{(B)} = \frac{\text{ChargingPotential}_s}{R_s^L + R_s^H} = \frac{\gamma \sum_{i \geq 1} \sum_{(i,j) \in \mathcal{K}_s} i \pi_s^{(\text{charged})}(i,j)}{R_s^L + R_s^H}, \quad (4.30)$$

where $\pi_s^{(\text{charged})}(i,j)$ represents the steady-state probabilities assuming the station is equipped with a charger. Normalizing by total expected outflow ensures that the score reflects the relative gain in high-SOC availability compared to the station's current usage. Stations with high scores are strong candidates for charger installation.

These scores collectively guide the selection of underperforming charger stations for removal and replace them with high-scoring non-chargers. After each candidate swaps and updates the charging status γ_s , the equilibrium computation is re-run to obtain a new total ridership $\mathcal{R}'$. If $\mathcal{R}'$ exceeds the previous $\mathcal{R}$ by a threshold, the swap is accepted; otherwise, the configuration is reverted. This iterative process continues until no further beneficial swaps are identified or until a pre-specified limit on the number of swaps is reached.

4.5 Analysis

4.5.1 Station-Level analysis of Markov Chain

We first analyze the two-dimensional Markov chain at station-level to gain deeper insights into how key parameters influence the chain. This preliminary analysis helps us understand the steady-state probabilities of different station states under various charging, rental, and arrival rate conditions.

We consider a station with a docking capacity of $C_s = 2$, which results in a total of 6 feasible states as $\{(0,0), (0,1), (0,2), (1,0), (1,1), (2,0)\}$ (see Figure 4.3-a). Each of these states corresponds to a steady-state probability $\pi_s(i,j)$, which represents the long-run probability of the station being in that state. In Figure 4.3-b to -d we illustrates how the steady-state probabilities

change in response to charging, low-SOC rental, and high-SOC arrival rate variations.




Figure 4.3: a) State transition diagram($C_s = 2$), station-Level analysis of steady state probabilities versus: b) Charging rate, c) Low-SOC rental rates, d) High-SOC arrival rate

As γ increases, Figure 4.3-b, states with a higher number of high-SOC bikes, such as $\pi(0,1)$ and $\pi(0,2)$, become more probable, while states dominated by low-SOC bikes, such as $\pi(1,0)$ and $\pi(2,0)$, decline. This trend is expected, as a higher charging rate mitigates the accumulation of low-SOC bikes. At sufficiently large values of γ , the probabilities of states with low-SOC bikes approach zero, indicating that nearly all low-SOC arrivals are converted to high-SOC almost immediately. As a result, further increases in γ yield diminishing returns. This highlights the importance of strategically placing charging infrastructure in stations with high low-SOC arrivals, ensuring an adequate supply of rentable high-SOC bikes without excessive investment in charging facilities. Figure 4.3-c indicates that as μ^L increases, the probability of low-SOC-dominant states declines, whereas states with fewer low-SOC bikes, such as $\pi(0,0)$ and $\pi(0,1)$, increase. However, increasing μ^L alone does not fully eliminate low-SOC accumulation, as the system remains constrained by arrival rates and charging availability. At higher μ^L values, the

system stabilizes, with diminishing effects on the steady-state probabilities, reinforcing the need for a balance between rental activity and charging efficiency.

The high-SOC arrival rate λ^H reduces reliance on charging infrastructure by increasing the proportion of bikes that arrive at a station already sufficiently charged (see Figure 4.3-d). As λ^H increases, states with a greater number of high-SOC bikes, such as $\pi(0, 1)$ and $\pi(0, 2)$, become more probable, while states dominated by low-SOC bikes decrease ($\pi(2, 0)$ and $\pi(1, 0)$). At high values of λ^H more bikes enter the system already charged. As a result, the station spends more time in states with fewer low-SOC bikes, making rental availability less dependent on charging. However, in stations that predominantly receive low-SOC arrivals (λ^L), charging capacity remains critical to ensure an adequate supply of high-SOC bikes for rentals. Across all variations, certain states exhibit non-monotonic behavior, initially increasing before declining. This pattern arises due to a redistribution of probability mass: as more high-SOC bikes become available, they are rented out at a higher rate, shifting probability among different states rather than following a strictly increasing or decreasing trend. These results underscore the interplay between bike rental dynamics, arrival patterns, and charging infrastructure.

To further investigate the station-level Markov chain dynamics, we analyze specific policy scenarios where station operations are constrained by limiting either arrivals or rentals. These scenarios provide insight into how charging rate variations interact with different demand patterns. These scenarios are defined as follows:

- 1) First scenario considers the case where no high-SOC bikes arrive at the station ($\lambda^H = 0$), meaning all incoming bikes are low-SOC. This situation reflects real-world stations where incoming bikes primarily arrive in a depleted state, such as at the endpoints of high-demand commuter routes. Without high-SOC arrivals, the system relies entirely on charging (γ) to maintain an adequate supply of high-SOC bikes, making this an important test for understanding the minimum charging requirements needed to sustain a functional station. Figure 4.4-a illustrates the state transition diagram of this scenario.
- 2) Second scenario investigates the effect of restricting low-SOC rentals ($\mu^L = 0$), where only high-SOC bikes can be rented (see Figure 4.4-b). This setup models user behavior in which riders avoid low-SOC bikes due to usability concerns specially for long-range trips. Examining

this case allows us to determine how restricting low-SOC rentals affects overall bike availability and whether charging alone can compensate for reduced turnover.

3) Third scenario combines these constraints, considering the case where all arrivals consist of low-SOC bikes ($\lambda^H = 0$) and low-SOC bikes cannot be rented ($\mu^L = 0$). This represents an extreme condition where high-SOC bikes can only be generated through charging, creating a worst-case scenario for bike availability (see Figure 4.4-c).

For each of these policies, we examine the steady-state probabilities in terms of varying γ to assess how the charging rate influences state distributions under constrained conditions. Results are presented in Figure 4.4-d to e.

In the first scenario ($\lambda^H = 0$), where all incoming bikes are low-SOC, low-SOC dominant states ($\pi(2,0)$, $\pi(1,0)$) initially have the highest probabilities. As γ increases, these states decline as more bikes transition to high-SOC, shifting probability toward mixed-SOC states. In the second scenario ($\mu^L = 0$), where low-SOC bikes cannot be rented, low-SOC bikes remain in the station until they are charged (since they cannot be rented out). At lower charging rates, low-SOC bikes may accumulate, increasing the likelihood of states such as $\pi(1,0)$ or $\pi(2,0)$. However as γ increases, low-SOC bikes are converted to high-SOC more quickly. Consequently, the system shifts toward states with more high-SOC bikes ($\pi(0,2)$ or $\pi(0,1)$). The third scenario ($\lambda^H = 0, \mu^L = 0$) represents the most constrained setting, where only charging generates high-SOC bikes. Low-SOC states are initially dominant, but as γ increases, high-SOC states gradually emerge (like $\pi(0,2)$ or $\pi(0,1)$). Since there are no direct arrivals of high-SOC bikes, this transition is slower than second scenario.

4.5.2 Computation Example with Exact Probability Model

To evaluate the performance of our proposed methodology and solution framework, we consider a computational example with a network of 20 bike-sharing stations, each having a capacity of three docks. Each station has predefined values for the rental rates and charging stations operate at a uniform charging rate of $\gamma = 0.6$ bikes per unit time. Initially, five stations are randomly selected to be equipped with charging infrastructure. The station interaction dynamics are governed by an origin-destination (O-D) distribution matrix, which quantifies the distribution of




Figure 4.4: Station-level Markov Chain analysis under different policies: (a–c) State transition diagrams, d) Steady-state probability vs. charging rate for first policy ($\lambda^H = 0$), e) Steady-state probability vs. charging rate for first policy ($\mu^L = 0$), f) Steady-state probability vs. charging rate for first policy ($\lambda^H = \mu^L = 0$).

trips between stations. This matrix defines the fraction of total trips that originate from one station and terminate at another, capturing the movement patterns within the network. Additionally, each station is assigned an SOC transition probability matrix, which defines the likelihood of e-bikes changing their SOC state while traveling between stations. Since on-route charging is not considered in this study, the probability of transitioning from low-SOC to high-SOC during transit is set to zero. Figure 4.5-a illustrates the spatial distribution of the stations and their corresponding O-D connections. The station locations are marked in blue, and the thickness of

the connecting edges represents the trip demand between stations. Darker and more prominent edges indicate higher interaction frequencies, while lighter edges signify lower demand.

We implement the exact probability model and compare the results of different heuristic search approaches as well as brute-force search to evaluate optimality. In ridership derivation of our solution methodology, the arrival rate convergence (λ^L and λ^H) at each station is tracked over iterations. Figure 4.5-b presents the evolution of these values, demonstrating how the system achieves equilibrium during the iterative process of the ridership derivation of the algorithm.




Figure 4.5: a) Station locations and O-D distribution, where edge thickness represents trip frequencies, b) Convergence of arrival rates across stations over iterations, confirming the attainment of equilibrium, c) Ridership trend across iterations for different approaches.

Table 4.1 presents a comparative analysis of the final results obtained from different approaches. The "Optimal Config." column denotes the indices of stations that were selected as optimal charging station locations and the underlined values indicate stations that also appear in the brute-force optimal configuration. The results indicate that the Direct Ridership Improvement (DRI) approach achieves a near-optimal solution with a ridership difference of only 0.04 compared to the brute-force approach, while reducing computation time dramatically from 1032.11s to 11.24s. Similarly, low-SOC Prevalence and Charging Potential approaches provide competi-

tive results with a slightly lower ridership difference of 0.13, but with significantly lower computational times (6.72s and 2.25s, respectively). These findings underscore the effectiveness of heuristic methods in achieving solutions that closely approximate optimal performance while improving computational efficiency. The trade-offs between ridership gains and computational cost suggest that heuristic approaches offer practical advantages, particularly in large-scale applications where brute-force methods become computationally prohibitive. In Sections 4.5.3, 4.5.4, and 4.5.5 we extend this analysis by exploring the sensitivity of these approaches under varying parameter configurations and to assess their robustness and scalability.

Figure 4.5-c illustrates the iterative improvement in ridership for three heuristic approaches compared to the optimal solution obtained via brute-force search. As the iterations progress, the Direct Ridership Improvement (blue) exhibits the fastest convergence, consistently outperforming the other heuristic methods and approaching the optimal ridership obtained by brute-force enumeration. This highlights its effectiveness in identifying near-optimal charger configurations with fewer iterations. Charging Potential (green) shows moderate improvement over iterations but lags behind the Direct Ridership Improvement approach. Low-SOC Prevalence (orange) demonstrates a slower rate of ridership enhancement, stabilizing at a lower final value compared to the other two heuristics. This figure confirms that all heuristic methods successfully improve ridership over time, but the rate of convergence and final performance vary.

Table 4.1: Comparison of Optimal Configurations, Ridership, and Computation Time.

Method	Optimal Config.	Ridership	Computation Time (s)
Direct Ridership Improvement	(7, 10, 13, 15, 18)	8.60	11.24
Low-SOC Prevalence	(5, 6, 10, 13, 15)	8.51	6.72
Charging Potential	(5, 6, 10, 13, 15)	8.51	2.25
Brute-Force	(4, 5, 7, 10, 15)	8.64	1032.11
Direct Ridership Improvement vs. Brute-Force		0.04	
Low-SOC Prevalence vs. Brute-Force		0.13	
Charging Potential vs. Brute-Force		0.13	

To assess the robustness of the proposed methodology, we perform a sensitivity analysis by varying key system parameters: the number of charging stations (N), the station capacity (C_s), and the charging rate (γ). This analysis offers insights into how increasing infrastructure capacity and charging capabilities influence e-bike availability, station occupancy, model performance and computational runtime. We consider a network consisting of twenty bike-sharing stations

with an initial setup of three charging stations and ten docking spots per station. We vary each parameter while keeping others fixed to evaluate their individual impact. In the subsequent analysis, we use Direct Ridership Improvement approach as the best performed approach. In the following, we will discuss the impact of each parameter on key performance metrics, as well as provide a comparison across them.

4.5.3 Impact of Number of Chargers

The number of charging stations plays a pivotal role in the planning and operation of e-bike sharing systems, directly influencing system performance, user experience, and financial feasibility. While increasing N is expected to enhance bike availability by reducing the prevalence of low-SOC bikes and improving overall ridership, it is essential to assess whether this increase continues to yield meaningful improvements beyond a certain threshold. Additionally, the computational cost of optimizing charger placement must be considered.

Figure 4.6-a illustrates how increasing N (the number of chargers) affects total ridership, comparing a heuristic-based Direct Ridership Improvement method with a brute-force approach as a benchmark for optimality. As expected, ridership generally increases with more chargers. However, the growth rate diminishes after a certain threshold, and the model even shows a slight decline at higher N . Although this decline remains relatively very small, it indicates that adding chargers past a certain threshold may not yield additional benefits and could introduce minor imbalances in bike distribution. In Markov chain equilibrium, even small changes in station-level probabilities (e.g., a slightly higher chance of being full at a critical station) can produce a small net loss in overall ridership. Because total ridership is summed across the entire network, localized negative shifts can outweigh improvements elsewhere, resulting in a modest decline despite many stations benefiting from additional chargers. Figure 4.6-b corroborates these findings by plotting the marginal ridership gain per additional charger. The gains initially spike when new chargers are introduced but drop sharply once the system's charging needs are largely met. Beyond this saturation point, the marginal benefit approaches zero or becomes very slightly negative, reflecting potential disruptions in station-level dynamics, such as certain locations becoming over-supplied with high-SOC bikes while others face deficits.




Figure 4.6: Number of chargers sensitivity analysis results. a) Total ridership and optimal ridership, b) Marginal ridership gain, c) Computation runtime, d) Fraction of high-SOC bikes for each station, e) Probability of being at capacity for each station, f) Optimal charger configurations.

Figure 4.6-c examines the computational runtime of the optimization approaches as N increases. The brute-force approach exhibits exponential growth in runtime, quickly becoming infeasible for large N . In contrast, the heuristic method scales much more efficiently, maintaining practical runtime performance even for large values. The divergence in computation time between the two methods is particularly relevant in real-world applications, where system planners must balance solution quality and computational feasibility. The heuristic approach enables near-optimal decisions without excessive computational overhead, making it a viable strategy for

large-scale, dynamic bike-sharing networks.

To gain further insights into the station-level impact of increasing N , Figure 4.6-d explores the fraction of high-SOC bikes at each station across different values of N . As more chargers are introduced, the availability of high-SOC bikes generally improves across the network. However, the effect is not uniform—some stations experience a substantial increase in high-SOC availability, while others show only marginal gains. This discrepancy suggests that the impact of charger placement is highly dependent on station rental rate, connectivity, demand patterns, and flows.

Figure 4.6-e complements this analysis by visualizing the probability of stations becoming full under different values of N . It shows that increasing number of chargers helps mitigate station overflows by improving the overall distribution of available bikes, thereby reducing system congestion and enhancing service reliability. Figure 4.6-f provides a heatmap of charger placement decisions across varying values of N , illustrating which stations consistently appear in the optimal configuration. This figure suggests that certain locations provide higher system-wide benefits, likely due to their connectivity or demand within the network.

4.5.4 Impact of Station Dock Capacity

The station dock capacity (C_s) is a fundamental parameter in e-bike sharing systems, as it determines the maximum number of bikes that can be accommodated at each station. Increasing C_s can reduce the likelihood of stations running out of available bikes, improving service reliability and ridership. However, expanding capacity indefinitely may lead to diminishing returns if demand saturation occurs or if bikes remain idle due to insufficient trip demand.

Figure 4.7-a illustrates how increasing C_s affects total ridership for both the heuristic-based Direct Ridership Improvement and the brute-force approach. As anticipated, raising C_s initially results in a significant increase in ridership, as more bikes are made available to meet demand. However, after reaching a certain threshold, the growth in ridership begins to slow down, indicating that adding more capacity no longer significantly enhances system performance. This trend is further explored in Figure 4.7-b, which showcases the marginal ridership gain associated with each incremental increase in C_s . The results demonstrate a sharp rise in ridership

gain when C_s is low, followed by a peak, after which the marginal benefits start to decline. Ultimately, increasing C_s has little effect on ridership, suggesting that an optimal capacity exists beyond which further expansion is unnecessary and has a negligible effect on ridership. Figure 4.7-c demonstrates that the brute-force approach exhibits a steep increase in runtime as C_s grows, making it impractical for larger dock capacities. In contrast, the heuristic method maintains significantly lower computational costs, confirming its effectiveness for large-scale implementations.




Figure 4.7: Dock capacity sensitivity analysis results. a) Total ridership and optimal ridership, b) Marginal ridership gain, c) Computation runtime, d) Fraction of high-SOC bikes for each station, e) Probability of being at full capacity for each station, f) Optimal charger configurations.

Figure 4.7-d explores the fraction of high-SoC bikes available at each station for different values of C_s . Same as increasing N in the previous section, The results show that the proportion of high-SoC bikes improves as dock capacity increases. However, the effect is not uniform across stations. A critical operational factor in station-based systems is the probability of stations becoming full, as full stations prevent new bike returns and disrupt system balance. Figure 4.7-e shows that increasing C_s significantly reduces the probability of stations reaching full capacity, particularly in high-demand areas. However, beyond a certain threshold, the probability of full stations approaches zero, indicating that further increases in C_s lead to underutilized capacity. Lastly, Figure 4.7-f reveals that certain stations consistently appear in the optimal charging configuration, regardless of C_s , while others are selected when dock capacity is high. This observation highlights the importance of jointly considering station capacity and charging infrastructure when designing an efficient system.

4.5.5 Impact of Charging Rate

The charging rate, γ , indicates how quickly bikes regain charge when docked at a charging-equipped station. This parameter directly impacts the availability of high-SOC bikes, which in turn affects ridership, station congestion, and overall system performance. However, the degree to which an increase in γ enhances these metrics depends on system constraints and demand dynamics.

Regarding ridership trends, Figure 4.8-a shows that increasing γ initially results in a significant rise in ridership, as anticipated. However, this trend quickly reaches a saturation point, demonstrating that beyond a certain threshold, additional increases in γ provide minimal benefits. This outcome suggests that excessively high charging speeds do not significantly enhance ridership, likely due to the influence of other limiting factors, such as station capacity. The marginal gains in ridership indicate that the most substantial increases come from the initial rise in the charging rate (Figure 4.8-b). At the same time, the benefits diminish sharply and approach zero as γ continues to increase. This finding highlights the importance of balancing operational efficiency with infrastructure investment when optimizing charging rates to avoid unnecessary costs with little improvement in ridership. The computational runtime of both the heuristic and brute-force




Figure 4.8: Charging rate sensitivity analysis results. a) Total ridership, b) Marginal ridership gain, c) Computation runtime, d) Fraction of high-SOC bikes for each station, e) Probability of being at full capacity for each station, f) Optimal charger configurations.

methods increases as γ rises (Figure 4.8-c); however, the heuristic approach maintains a more stable computational cost across different γ values.

The effect of γ on the availability of high-SOC bikes is illustrated in Figure 4.8-d. As expected, increasing γ generally enhances the proportion of high-SOC bikes at stations, particularly at locations equipped with chargers. However, this effect plateaus beyond a certain threshold, indicating that further increases in charging speed yield minimal additional benefits. This suggests that other system constraints, such as station dock capacity, become the dominant limiting factors beyond a certain point. The heatmap in Figure 4.8-e shows the probability of stations

reaching full capacity under different γ values. The results suggest that higher charging rates do not significantly increase the risk of station overflows. While localized congestion may occur, the overall impact of γ on system-wide availability remains controlled. Although charging speed can enhance bike availability, it does not inherently introduce major system-wide bottlenecks. Figure 4.8-f highlights the charger placement patterns across different γ values. Certain stations, specifically stations 4, 10, and 15, consistently appear in the optimal charging configuration across a wide range of charging rates. This finding suggests that these stations play a crucial role in network efficiency, likely due to their strategic location within the system. Additionally, increasing γ only affects stations that were already designated as charging locations in the given configuration, reinforcing the idea that the placement of chargers is more influential than the charging rate itself. This highlights the importance of optimizing station selection rather than simply increasing charging speed indiscriminately.

4.5.6 Comparative Impact of Key System Parameters

As the final part of the sensitivity analysis, we compare the effects of station dock capacity (C_s), the number of chargers (N), and the charging rate (γ) across key performance metrics, including ridership, marginal ridership gain, and computational runtime. The results reveal clear trends that highlight the importance of each parameter in influencing system performance.

Among the three factors, increasing station dock capacity (C_s) has the most significant impact on ridership, leading to substantial growth. Figure 4.9-a indicates that system capacity is often the primary constraint on achieving ridership goals. While the number of chargers (N) also positively affects ridership, it demonstrates diminishing returns; additional chargers become less effective after reaching a certain threshold. In contrast, increasing the charging rate (γ) has a limited effect, with ridership gains stabilizing following an initial increase. These diminishing returns are further illustrated in the marginal ridership gains presented in Figure 4.9-b. Here, dock capacity initially offers the highest benefits, but the gains quickly taper off. The trend for additional chargers is similar, while increasing γ leads to a notable decline in marginal gains, reinforcing the idea that a moderate charging rate is adequate for optimal system performance. From a computational standpoint, expanding dock capacity significantly increases runtime, re-




Figure 4.9: Comparison of changes in key system parameters (number of chargers, dock capacity, charging rate) on: (a) Total ridership, (b) Marginal ridership gain, and (c) Computation time.

flecting the complexity of handling larger state spaces (see Figure 4.9-c). While the number of chargers also affects computation time, this impact is more manageable. In contrast, variations in γ have minimal influence, indicating that changes in charging rates do not substantially increase computational overhead. These findings highlight the necessity for a balanced infrastructure strategy. Although expanding dock capacity and optimizing charger placement are essential for enhancing ridership, their diminishing returns underscore the importance of identifying thresholds. While adjustments to charging rates can be beneficial at first, they tend to plateau quickly, making the capacity of stations and the allocation of chargers more critical factors for long-term system efficiency.

4.5.7 Comparison of Exact Probability vs. Single-Pooling State Approximation Model

The necessity of using the single-pooling approximation arises from the computational challenges observed in sensitivity analysis results. While the exact probability model provides high accuracy, its computational cost increases significantly with system size, particularly as the station capacity (C_s) grows. As a result, an alternative approach is needed to ensure computational feasibility while maintaining reasonable accuracy. In this section, we compare the performance of the exact probability model and the single-pooling state approximation, considering the same network configuration across different station capacities (C). In both models we use the Direct Ridership Improvement approach, the best performed one, in the heuristic stage.

Table 4.2 presents the comparative results for both the exact probability model and the single-

pooling state approximation model, benchmarked against the brute-force approach. The comparison is conducted across various station capacities ($C = 3, 4, 10, 15, 30$), evaluating their ability to identify the optimal charger placement configuration, the resulting system ridership, and the computational cost.

Table 4.2: Comparison of exact probability vs. single-pooling state approximation model

Capacity	$C = 3$	$C = 4$	$C = 10$	$C = 15$	$C = 30$
Exact Probability Model					
Optimal Config.	(0, 2, 5)	(0, 5, 6)	(0, 2, 5)	(0, 2, 9)	(0, 2, 9)
Ridership	4.28	6.22	19.67	29.96	30.00
Computation Time (s)	1.21	2.62	17.43	673.10	4086.10
Single-Pooling State Approximation Model					
Optimal Config.	(0, 2, 5)	(0, 5, 6)	(0, 2, 5)	(0, 2, 9)	(0, 2, 9)
Ridership	4.06	6.18	19.29	29.55	30.00
Computation Time (s)	4.10	3.47	27.57	38.03	77.24
Brute-Force Search					
Optimal Config.	(0, 5, 6)	(0, 5, 6)	(0, 2, 5)	(0, 6, 9)	(0, 6, 9)
Ridership	4.28	6.22	19.67	29.97	30.01
Computation Time (s)	3.20	5.23	52.42	1830.02	16271.57

As seen in Table 4.2, the primary distinction between the models lies in their computational efficiency. As expected, the brute-force approach exhibits an exponential increase in computation time as station capacity grows. For instance, at $C = 15$, the brute-force method requires 1830.02 seconds, whereas the exact probability model completes the optimization in 673.10 seconds, achieving a notable reduction. However, as C further increases, the exact probability model also experiences a substantial computational burden, demanding 4086.10 seconds at $C = 30$, making it increasingly impractical for large-scale networks. Conversely, the single-pooling state approximation model significantly reduces computation time across all scenarios. Notably, when $C = 30$, it completes 53 times faster than the exact probability model and over 200 times faster than brute-force enumeration. This drastic improvement in efficiency is illustrated in Figure 4.10-a and underscores the effectiveness of the approximation method for large-scale applications where computational feasibility is a key concern. Both models effectively identify the optimal charger configuration in most cases, closely aligning with the brute-force results. However, minor deviations occur at certain capacities. Despite these small discrepancies, the ridership values obtained remain very close to those determined by brute-force enumeration,




Figure 4.10: Models comparison in terms of station capacity and a) Computation time, b) Ridership

demonstrating the high accuracy of both models in capturing system equilibrium behavior. Figure 4.10-b highlights the minimal differences in ridership values between the methods. The almost identical ridership trends across varying station capacities confirm that the single-pooling approximation remains a highly effective heuristic, successfully capturing the key system dynamics despite its simplified state-space representation.

In addition to evaluating ridership and computational efficiency, we further analyze the steady-state probability distributions obtained from both models. This analysis assesses the accuracy of the single-pooling state approximation model relative to the exact probability model in capturing the stationary behavior of the system. Figure 4.11 presents a scatter plot comparing the steady-state probabilities computed by the single-pooling approximation model against those obtained from the exact probability model for a network of 10 stations. Each point represents a probability value for a given station-state pair, while the dashed red line represents a perfect one-to-one correspondence between the two models.

The results indicate that the single-pooling approximation closely follows the exact model with most points aligning well along the reference line. However, minor deviations can be observed, particularly for some stations where the approximation slightly overestimates or underestimates the state probabilities. These differences are expected due to the simplifications inherent in the single-pooling model, which aggregates system states into a single representative state and then distributes the probability based on the binomial approximation. To quantify these differences, we compute the Total Variation Distance (TVD) for each station and the overall average TVD across all stations, which is 0.052448. This low TVD value demonstrates that the overall dis-




Figure 4.11: Comparison of steady-state probabilities in single-pooling approximation model and exact probability model for a network of 10 stations.

crepancy between the two models is minimal, reinforcing the validity of the single-pooling state approximation.

These findings demonstrate that the single-pooling state model offers a practical balance between computational efficiency and accuracy, making it suitable for large-scale and high capacity implementations where exact solutions become computationally prohibitive.

4.6 Real Case Studies: Pittsburgh, Vancouver, San Francisco

We apply our model to three cities with varying bike-sharing system sizes, using publicly available data. Pittsburgh, US, operated by POGO, represents a smaller system; Vancouver, Canada, operated by Mobi, a mid-sized system; and San Francisco, US, operated by Lyft, a large-scale system. For each city, we use the most recent ridership data, including departure and return timestamps, station names, and trip distances or durations; if available. This data is enriched using station-level information from the General Bikeshare Feed Specification (GBFS), including station names, geographic coordinates, and capacities. The integration of these sources supports in-depth analysis of mobility patterns and enables model preprocessing and parameterization.

A key preprocessing step involved refining the data to reflect permanent infrastructure and typical user behavior. We assumed a percentage of overall ridership to be electric, given varying

availability of explicit e-bike usage data. We excluded records with missing data or those associated with temporary or unlisted stations, removing trips related to bike testing, maintenance, and other non-user activities. This filtering reduced noise and improved model fidelity. After preprocessing, the final datasets included 58 stations in Pittsburgh, 257 in Vancouver, and 331 in San Francisco. From this data, we derived key model parameters. The origin-destination (O-D) matrix was constructed by aggregating trips between station pairs. Station-specific arrival (λ_s) and departure (μ_s) rates were estimated based on ridership data and trip frequencies. As direct SOC data is unavailable, we approximated the split between high- and low-SOC arrivals. Arrival rates (λ_s^H, λ_s^L) were estimated by distributing incoming trips based on SOC depletion inferred from trip distance. SOC transition probabilities (P_s^{XY}) were calibrated using a distance-dependent approach: trips with longer distances were more likely to deplete the battery, while shorter trips preserved high SOC levels. These estimated parameters were then incorporated into the algorithm.

In each city, we conducted a sensitivity analysis regarding the number of chargers to evaluate the model's performance across different levels of electrification. In Pittsburgh's smaller, more compact system, we examined scenarios with approximately 10%, 30%, and 50% station electrification, capturing progressive expansions in charger coverage (Figure 4.12). In the initial scenario (Figure 4.12-a), the algorithm strategically selects a limited number of high-throughput stations, such as those near Downtown. As charger availability increases (Figures 4.12-b and 4.12-c), coverage extends into additional high-demand areas, enhancing accessibility and reducing low-SOC occurrences.




Figure 4.12: Pittsburgh bike-sharing network: Optimal charging-station distributions for approximately 10, 30, 50 percent electrification.

For Vancouver and San Francisco, we focused on two electrification scenarios—approximately

10% and 50%—to illustrate the impacts of sparse versus near-saturated charger coverage in mid-sized and large-scale systems (Figure 4.13). In Vancouver, Figures 4.13-a and 4.13-b depict optimal charger placements for 10% and 50% electrification, respectively. Similarly, for San Francisco, Figures 4.13-c and 4.13-d illustrate charger distributions.

In these larger cities, SOC probability transitions significantly impact charger optimization due to greater variability in trip distances and usage patterns. Longer trips driven by sprawling urban layouts and diverse user behaviors increase the likelihood of battery depletion, necessitating a more robust charging infrastructure. Our heuristic algorithm iteratively optimizes charger locations by integrating SOC dynamics with ridership patterns, prioritizing stations that maximize charging potential and minimize downtime. Understanding these dynamics could further inform the scalability of charging networks as e-bike systems grow and adapt to shifting urban mobility demands. This ensures that the model scales effectively across cities of varying sizes, adapting to each network’s unique spatial and operational characteristics.




Figure 4.13: Vancouver and San Francisco bike-sharing networks: (a) and (b) optimal charging-station configurations in Vancouver for approximately 10% and 50% station electrification; (c) and (d) optimal charging-station configurations in San Francisco for approximately 10% and 50% station electrification.

4.7 Conclusion

This study presents a framework for optimizing charger placement in station-based e-bike sharing systems, addressing the critical interplay between SOC dynamics and ridership maximization. By modeling the system as a two-dimensional Markov chain, we capture the birth-death processes of low- and high-SOC bikes and charging transitions, providing a robust representation of the operational dynamics inherent in e-bike networks. The charger placement problem is formulated as a station-level location optimization model that maximizes system-wide ridership under station capacity and SOC constraints. To tackle the computational challenges of this high-dimensional, nonlinear optimization problem, we propose a local-improvement heuristic algorithm that iteratively refines charger configurations using multiple search strategies—Direct Ridership Improvement, Low-SOC Prevalence, and Charging Potential. Additionally, we introduce simplification methods, such as the single-pooling state approximation, to ensure scalability for large-scale networks where exact solutions are computationally prohibitive.

Our results demonstrate the efficacy of the proposed framework in achieving near-optimal charger configurations with reduced computational overhead. The heuristic algorithm consistently enhances ridership across various system configurations, with the Direct Ridership Improvement approach showing the fastest convergence and highest final performance. Sensitivity analyses reveal key insights into the impact of system parameters on performance: increasing the number of chargers generally improves high-SOC bike availability and reduces station congestion, but the benefits diminish beyond a certain threshold, and the effect varies across stations due to differences in demand, connectivity, and flows. Station-level analyses further highlight the importance of strategic charger placement—certain stations, often those with high demand or network centrality, are consistently prioritized in optimal configurations, underscoring their system-wide impact. The single-pooling state approximation proves to be a practical trade-off, maintaining reasonable accuracy while drastically reducing computational complexity, making the framework viable for real-world deployment in large, high-capacity networks.

The real-world case studies in Pittsburgh, Vancouver, and San Francisco validate the framework’s applicability, demonstrating its ability to adapt to diverse urban contexts and operational

constraints. These findings have implications for urban mobility planners and e-bike sharing operators. By providing a data-driven approach to charger placement, the framework supports the design of efficient, reliable, and sustainable e-bike systems that can meet growing demand. This study opens several avenues for future research. First, incorporating dynamic demand patterns and time-varying rental rates could enhance the model's realism, capturing the temporal fluctuations inherent in urban mobility systems. Second, extending the framework to account for multi-modal interactions—such as integration with public transit or other micromobility options—could further improve its utility in supporting transportation planning. Third, exploring the impact of user behavior, such as preferences for high-SOC bikes or incentives for returning bikes to charger-equipped stations, could provide deeper insights into system performance and user satisfaction.

Chapter 5

Decision-Support Tool Development and Practical Implementation

5.1 Overview and Motivation

The preceding chapters introduced a suite of analytical models designed to enhance bike-sharing systems' design, expansion, and electrification. These models include predictive, optimization, and infrastructure planning frameworks. While these models offer valuable methodological advancements, their full potential lies in their ability to support real-world decision-making. Without practical tools for planners and policymakers to engage with these models, the insights remain largely theoretical. This chapter addresses this challenge by presenting **BikesharePro**, a web-based, GIS-enabled decision-support tool that brings to life the research outputs developed throughout this dissertation.

Bike-share planning traditionally relies on static master plans that are developed every four to five years by third-party consultants. These plans are based on long-term projections and broad strategic visions. While they often consider important factors like equity, accessibility, and integration with public transit, they lack the flexibility needed to adapt to rapid changes. As cities experience shifts in land use, demographics, and travel behavior, these static plans can become outdated. This can result in poor station placement, inefficient use of bicycle fleets, underutilized stations, gaps in accessibility, inequitable service distribution, and a reduced ability to respond to emerging transportation needs.

To address these challenges, this dissertation introduces *BikesharePro*, an interactive, GIS-enabled planning tool that connects analytical modeling with practical implementation. y introducing a model-driven tool, we can facilitate a shift towards continuous planning, an approach

that enables cities to iteratively update, test, and adapt bike-share networks based on changing data and policy objectives. The tool combines descriptive analytics (e.g., ridership patterns, station usage, and socio-demographic overlays), predictive analytics (e.g., network-based demand forecasting), and prescriptive analytics (e.g., optimization algorithms for station siting). Through a responsive, interactive interface, planners and policy analysts can visualize existing conditions, simulate "what-if" scenarios, and optimize infrastructure investments effectively. This chapter proceeds as follows. First, we describe the architectural framework of Bikeshare-Pro, detailing its technical design and modular structure. Next, we outline the data sources and processing workflows that ensure consistency and usability across diverse contexts. We then present the tool's key functionalities, organized into panels for filtering, editing, analysis, optimization, and reporting, and describe how they operationalize the underlying models. We present real-world examples to illustrate the tool's scalability, adaptability, and policy relevance. The chapter concludes with a case study of Toronto that demonstrates the tool's applicability, flexibility, and policy relevance. We demonstrate how this tool advances analytical rigor and delivers actionable, user-friendly solutions for sustainable, equitable, and demand-responsive bike-share planning.

5.2 System Architecture and Data Integration

The proposed tool is a modular, web-based decision-support platform that combines data ingestion, modeling, visualization, and optimization into a single interface. It was developed using JavaScript and Mapbox for the front-end, incorporating optimization and prediction modules inspired by models discussed in earlier chapters. The architecture is designed to facilitate continuous planning and rapid deployment by using scalable data sources, automated processing, and dynamic user interaction. The system integrates multiple data layers obtained from open data portals, APIs, and mobility service providers. These layers are standardized through a common geospatial framework, ensuring consistency with the analytical models developed in this dissertation. The key data inputs include:

- **Station Properties:** Data on station capacity, year of deployment, e-bike and charger

availability, and fleet distribution, typically accessed from GBFS feeds or open data portals.

- **Transit Data:** Location of nearby transit stops, obtained through GTFS feeds, supporting integration analysis for first- and last-mile planning.
- **Socio-demographic Indicators:** Census-based metrics such as income, population density, age distribution, and commuting behavior, enabling descriptive analysis, demand prediction and equity analysis.
- **Ridership Transactions:** Origin-destination flows and pickup/drop-off timestamps between station pairs, derived from historical data published by bike-share operators or municipal data platforms.
- **Points of Interest (POIs):** Location-based features including schools, parks, business districts, and recreation areas, sourced via APIs to assess land-use contexts and attraction potential.
- **Equity Metrics:** Custom equity indicators, such as Neighbourhood Improvement Area profiles or access to essential services, either provided by cities or generated through synthesized indicators.

To effectively align and analyze these inputs, the tool uses a hexagonal grid system with cells typically ranging from 100 to 200 meters in radius, which corresponds to walkable access distances. These geospatial units serve as a scalable basis for aggregating features, allowing for proximity calculations, clustering of nearby stations, and overlaying of contextual data. This structure adheres to operational constraints related to station spacing, infrastructure availability, and multimodal connectivity. The internal data pipeline follows the preprocessing strategies established in previous chapters, ensuring consistency between the tool's outputs and the underlying analytical models. This alignment of data structures, modeling assumptions, and spatial resolution enables BikesharePro to function as a practical extension of the methodological contributions presented in the dissertation.

A conceptual overview of the system architecture is illustrated in Figure 5.1. This diagram outlines the flow from raw data to user-facing functionality, highlighting modularity and responsiveness as key design principles. It includes aspects such as data ingestion, preprocessing, modeling, analytics, and deployment.




Figure 5.1: System architecture of BikesharePro: from data sources to deployment.

5.3 Core Functionalities and Design

Figure 5.2 illustrates a snapshot of the BikesharePro tool calibrated for Toronto. In this figure, orange hexagons represent existing bike-sharing stations, while light orange hexagons denote their respective service areas. Blue hexagons depict areas without a station, color-coded by underlying spatial indicators such as age demographics, income levels, transit accessibility, or equity metrics. This enables a contextualized and data-driven understanding of unmet demand

and service equity.

The tool also overlays spatial context with bivariate flow visualizations at station locations: green circles represent station-specific outflows (trips originating from the station), while red circles represent inflows (trips ending at the station). The radius of each circle is proportional to the total volume of flows, visually highlighting demand intensity and imbalance at each station. Both circles are placed at the exact locations, allowing for a quick comparative assessment of inflow versus outflow.

The tool is organized into modular interface panels, each designed for specific planning activities and analysis categories, ranging from descriptive analysis (data exploration) to predictive and prescriptive analysis for infrastructure optimization. In the following, we describe the functionality and purpose of each panel in detail.




Figure 5.2: Snapshot of the tool calibrated for Toronto. Hexagon fills represent spatial coverage and station presence, while circle overlays depict station-level outflow (green) and inflow (red), scaled by volume. The underlying blue hex grid encodes demographic and accessibility indicators.

The Map View and Filtering Panel enables users to explore live maps with multiple overlay layers, including demographic indicators, station typologies, cycling infrastructure, and ridership over time. Filtering capabilities support targeted queries based on deployment year, station characteristics, ridership volume, transit integration level, and cycling congestion in cycle-friendly roads. These descriptive functionalities allow for detailed analyses of network growth, route categories, and modal connectivity. Built-in charts and spatial visualizations further support the

assessment of station-level indicators such as inflow/outflow volumes, utilization rates, imbalances, proximity to POIs, and financial sustainability. Figure 5.3 illustrates examples of these visual components.




Figure 5.3: Descriptive analytics capabilities example of the tool demonstrated through selected panels. a) OD ridership shown as volume-colored lines, filterable by year and season, b) Transit integration heatmap: red areas indicate no bike-share station within a walkable distance of a transit stop, yellow areas indicate one station, and blue areas indicate two or more, c) Cycling congestion on bike-friendly routes based on ridership and routing API, d) Station-level view with points of interest and monthly bike availability, e) Net flow visualization at station level with shaded circles—green for net outflow, red for net inflow—and corresponding system-level ridership metrics in the lower panel, f) Cost and revenue inputs with financial summaries for system modifications.

To support detailed scenario refinement, the tool provides Cell and Station Property Editors that enable planners to examine and modify the local attributes of each grid cell and station

(Figure 5.4). Users can directly modify station characteristics, such as dock capacity or charger availability, or add proposed stations to simulate expansion scenarios. These edits trigger recalculations across the analysis and prediction engine.




Figure 5.4: Interactive data editing panels. a) Cell-level demographics and land use editor, b) Station property editor with capacity, charger availability, and status settings.

An integral component of the tool is its predictive analytics engine, which inspired from the Graph Neural Network model introduced in Chapter 2. As users introduce new stations or modify existing ones, the tool dynamically predicts station-to-station ridership flows, leveraging station attributes, spatial adjacency, and network topology. These OD predictions are visualized on the map through directional flow lines and circle plots, enabling planners to assess demand sensitivity and connectivity impacts (Figure 5.5).




Figure 5.5: Predictive analytics for station deployment. a) Adding a new station with configurable attributes such as capacity, installation year, and charging capability, b) Predicted OD flows from the newly added station visualized as directional lines.

The prescriptive module aids infrastructure investment by recommending optimal locations for

new stations. It uses a multi-objective optimization algorithm that considers weighted planning goals, including accessibility, equity, and transit integration, while also prioritizing ridership potential, which is often overlooked in bike-sharing master plans. This optimization approach builds on the force-based algorithm introduced in Chapter 3, aligning station positions with demand hotspots and ensuring optimal station spacing (see Figure 5.6).




Figure 5.6: Prescriptive station placement. Both figures use the same weights for accessibility, transit integration, and equity. Red hexagons indicate newly recommended station locations. a) With low connectivity importance, the model favors outer, underserved regions, b) With higher connectivity importance, the model prioritizes integration with the existing network, placing stations closer to current infrastructure.

To support policy evaluation, the tool enables planners to generate side-by-side comparisons between baseline conditions and optimized configurations, facilitating scenario analysis and reporting. Numerical summaries and visual outputs clearly highlight changes in key metrics such as system coverage, utilization, and equity distribution. This tool is designed as a full-stack web application, allowing users to save and reload network configurations for iterative design exploration. These features support longitudinal comparisons, collaborative scenario reviews, and integration with broader planning workflows. Additionally, transportation planners can export summary tables and geospatial shapefiles to facilitate reporting, documentation, and integration with external GIS tools.

These integrated capabilities establish BikesharePro as a comprehensive planning environment which is beyond a model viewer and enables descriptive, predictive, and prescriptive analysis to inform tactical decision-making in varied urban contexts.

5.4 Scalability and Contributions

The proposed tool has been developed to accommodate a range of urban environments and planning needs, demonstrating strong scalability and flexibility across diverse city contexts. It has been successfully implemented in many cities, such as Toronto, Vancouver, Chicago, Boston, Pittsburgh, and Cincinnati (Figure 5.7). Additionally, BikesharePro provides a viable solution through transfer learning and model generalization for cities with limited historical ridership data, such as Ottawa. The tool's underlying predictive models, trained on normalized datasets from the aforementioned cities, enable the estimation of station-to-station ridership in new contexts based on structural, demographic, and spatial similarity measures. This capability extends the tool's utility to data-sparse regions, supporting proactive planning and equitable infrastructure deployment from the early stages of network development.




Figure 5.7: Deployment of BikesharePro across diverse urban environments. It has been applied using real operational data in Vancouver, Boston, San Francisco, Chicago, and Cincinnati, and evaluated with proxy data and similar models in Hamilton, Ottawa, and Mississauga.

In addition to its demonstrated versatility, BikesharePro introduces methodological innovations that reshape how cities approach strategic bike-share planning. Rather than relying on static master plans updated every four to five years, the tool facilitates continuous, responsive planning

through its integrated platform. This interactive planning interface enables user-driven scenario testing, live data manipulation, and strategic exploration of trade-offs. By using the predictive and optimization frameworks outlined throughout this dissertation, BikesharePro provides a practical, forward-thinking tool for data-informed urban mobility planning. It democratizes access to complex models and helps bridge the gap between theoretical research and actionable policy, ultimately contributing to the development of more resilient, equitable, and effective bike-sharing systems.

The following case study demonstrates how the tool can guide infrastructure decisions, test policy-driven scenarios, and generate actionable insights within a complex urban bike-sharing system.

5.5 Case Study: Toronto

This section highlights the real-world applicability of the developed tool by focusing on Toronto's bike-sharing system. As one of North America's largest programs, Toronto's system operates within a complex urban context marked by diverse population densities, varied multimodal travel needs, and an increasing emphasis on equitable access and electrification. The case study demonstrates how the tool can support scenario-based planning and empower decision-makers to balance multiple, often competing, planning goals.

We consider four key planning questions for shared docked bike-sharing systems:

- Which areas present the greatest potential for adding new stations to maximize overall ridership?
- How can we position new station locations in alignment with other transportation modes, such as public transit, to enhance first- and last-mile connectivity?
- What strategies can be implemented to expand access to underserved neighborhoods while ensuring operational feasibility and consistent usage?
- What is the potential impact on ridership and finances if we transition part of the fleet to e-bikes, including the necessary charging infrastructure?

5.5.1 Scenario Design and Objectives:

To address the four key planning questions, we construct a series of forward-looking scenarios using BikesharePro’s optimization and analytics modules. These scenarios reflect distinct priorities commonly faced by urban mobility planners, including ridership growth, multimodal integration, equity, and electrification. We begin with a baseline scenario representing the current state of Toronto’s bike-sharing system, followed by four alternative planning scenarios:

Scenario 0- Baseline: This scenario captures the most recent condition of the system, which includes 839 stations (23 of which are equipped with charging infrastructure) and serves approximately 3.67 million annual trips. The data is based on cleaned and filtered information from publicly available sources, with records containing missing or inconsistent trip or station information removed to ensure accuracy in scenario modeling. This baseline serves as a reference point for evaluating the performance and impact of proposed system changes.

Scenario 1- Ridership-focused expansion: This scenario addresses the first planning question by identifying areas with high potential for ridership growth. Using the tool’s weighted station suggestion module, new stations are found in zones characterized by high population density and proximity to major POIs (e.g., commercial hubs, educational institutions, and event venues). These locations are likely to generate high travel demand due to their role as consistent trip attractors and generators.

Scenario 2 – Transit-integration: This scenario aims to enhance first- and last-mile connectivity by placing new stations near major transit-rich areas. The station suggestion module is configured to prioritize transit integration as the main objective, while also considering population density and equity.

Scenario 3 – Equity-prioritized: This scenario introduces new stations in underserved areas, guided primarily by the City of Toronto’s Neighbourhood Improvement Area (NIA) designations. The station suggestion module is used with a weighting configuration that emphasizes equity, while demographic overlays (such as income, employment, and age) serve as supplementary indicators to refine station placement. This scenario aims to enhance access and coverage while maintaining feasible operational investments.

Scenario 4 – Electrified Ridership Expansion: This scenario builds on the ridership-focused

expansion plan (Scenario 1) by equipping 50% of the newly added stations with charging infrastructure to support the deployment of e-bikes. Among the new stations, those selected for electrification are prioritized based on their projected outflow volume, ensuring that high-demand origin points are equipped to maintain reliable availability of charged bikes. These locations, situated in dense, trip-generating areas near key points of interest, are particularly well-suited for e-bike usage due to longer trip distances and elevated service pressure.

5.5.2 Optimization and Analytical Workflow

We generate candidate station locations using the *Prescriptive Optimization Module*, which identifies high-impact locations based on the weighted planning priorities for each scenario (Table 5.1). The resulting priority distributions and corresponding optimized station locations for Scenarios 1–3 are shown in Figure 5.8. The spatial patterns of Scenario 1 and 4 are similar due to shared weighting schemes; however, the addition of charging infrastructure in Scenario 4 alters downstream performance metrics, including capital costs, system revenue, and station-level usage.

Table 5.1: Summary of Planning Scenarios

Scenario	# New	Access.	Transit.	Equity	Connectivity	% Charging
Ridership-focused	30	60%	20%	20%	0.8	0%
Transit Integration	30	20%	70%	10%	0.5	0%
Equity-Prioritized	30	10%	30%	60%	0.5	0%
Electrified Ridership-focused	30	60%	20%	20%	0.8	50%

It is important to note that these configurations are not presented as final or static recommendations. Rather, they reflect one iteration within an exploratory planning environment. The tool allows planners to adjust inputs, test alternative assumptions, and generate new configurations on demand.

As shown in Figure 5.8, each scenario generates distinct spatial priorities and optimized station placements. The ridership-focused expansion concentrates on dense, trip-generating areas in downtown and midtown Toronto and the transit-integration scenario shifts attention toward corridors and nodes with high multimodal potential. In contrast, the equity-prioritized configuration expands service into underserved neighborhoods. These spatial patterns reflect how the




Figure 5.8: Scenario-specific spatial weighting distributions and corresponding optimized station locations. Figures (a), (c), and (e) display the weighted planning maps for Scenario 1 (Ridership-focused), Scenario 2 (Transit Integration), and Scenario 3 (Equity-Prioritized), respectively. Darker blue shading indicates higher priority based on the weightings. Figures (b), (d), and (f) show the resulting optimized new station locations, highlighted in red, as generated by the Prescriptive Optimization Module.

tool translates planning goals into differentiated infrastructure recommendations.

Once candidate station locations are generated, the *Descriptive Analytics Module* is used to evaluate their spatial context, considering factors such as proximity to POIs, demographic characteristics, and network topology. These insights inform the subsequent application of the *Pre-*

dictive Ridership Module, which estimates demand impacts after incorporating the proposed stations into the network. In select cases, locations are manually adjusted to adjacent cells based on planner expertise or local constraints. This hybrid approach underscores the tool’s role as a flexible decision-support system that integrates data-driven modeling with spatial responsiveness and practical judgment.

5.5.3 Results and Comparative Insights

We assume a phased implementation schedule, with 50% of the stations deployed in 2025 and the remainder in 2026. Scenario performance is evaluated across a comprehensive set of metrics encompassing system-level and spatial outcomes. These include changes in aggregate ridership, system revenue, infrastructure investment requirements, and network and operational indicators such as net flow balance, congestion on cycling-friendly corridors, alignment with major transit infrastructure, and service coverage in underserved neighborhoods. Summary results are provided in Table 5.2.

Table 5.2: Comparative Performance Metrics Across Planning Scenarios

Performance Metric	Baseline	Ridership-focused	Transit Integration	Equity-Prioritized	Electrified Ridership
Projected Ridership	3.67M (-)	3.79M (+3.27%)	3.78M (+3.0%)	3.75M (+2.2%)	3.81M (+3.81%)
Transit Integration Score	173	174	176	177	173
% of Transit Zones Upgraded	–	0.15%	0.49%	0.48%	0.15%
Estimated Revenue (\$)	–	+1.21M	+1.08M	+0.80M	+1.31M
Net Flow Imbalance Index (NFII)	34.17	34.86	35.06	34.27	34.08

To assess the extent to which each scenario enhances multimodal connectivity, we report a “Transit Integration Score”, which quantifies the alignment between the bike-share network and major transit infrastructure. Higher scores indicate stronger integration, reflecting improved support for first- and last-mile connectivity. In addition to the aggregate score, we track the per-

centage of transit-rich zones that experienced upgrades in network accessibility, specifically those that moved from low (red) to moderate (yellow) or high (blue) access following the addition of new stations (Refer to Figure 5.3-b).

We also assess system balance using a "Net Flow Imbalance Index (NFII)". This metric quantifies the average absolute difference between inflow and outflow at each station, providing insight into how well the system supports continuous service without requiring excessive fleet relocation. A lower NFII indicates a more self-balancing network, where stations neither frequently empty out nor overflow, thereby reducing unmet demand and operational burden.

5.5.4 Concluding Reflections, Value, and Future Directions

The comparative analysis highlights trade-offs between ridership, equity, and multimodal goals. Ridership-focused scenarios maximize usage and revenue but offer limited improvements in access or transit integration. In contrast, equity-oriented plans enhance coverage in underserved areas, though with lower demand. Transit-integrated scenarios strike a balance between connectivity and performance. These differences demonstrate the value of scenario-based planning, and underscore BikesharePro's utility as a flexible decision-support tool that translates evolving policy goals into adaptive, data-driven infrastructure strategies.

At the same time, some limitations must be acknowledged. While the predictive ridership module incorporates spatial network effects, it assumes behavioral stability and does not yet capture dynamic responses such as induced demand, evolving mode preferences, or feedback effects from policy interventions or service adjustments. These limitations are typical of strategic planning tools and are expected to diminish as richer, more temporally granular data sources become available. Additionally, although the station placement logic leverages spatial demand patterns and demographic overlays, it remains bound by the resolution, accuracy, availability of input datasets, and city-specific geographic constraints. Lastly, while the tool provides high-level insights into station-level imbalances and supports relocation-aware planning, operational redistribution decisions remain outside its intended scope.

Future enhancements will focus on expanding analytical capabilities and practical utility. As more datasets become available, predictive models will incorporate temporal variation, be-

havioral feedback, and induced demand to improve forecast accuracy. Optimization features will support advanced electrification strategies, including charger deployment at constrained or legacy stations, with prioritization based on grid capacity, topography, and route gradients. The cost and revenue module will be refined to model operating costs, subsidies, and peak demand more granularly. Transit integration features will include joint optimization of bike-share and transit nodes, transfer penalties, and multimodal trip chaining.

Chapter 6

Conclusions and Future Work

6.1 Summary of the Research

This dissertation developed an integrated, data-driven framework for planning station-based bike-sharing systems, addressing three core challenges: (i) predicting OD ridership by capturing spatial and structural dependencies within the network; (ii) optimizing station locations to maximize ridership rather than just coverage; and (iii) integrating e-bike infrastructure planning through models that capture charging constraints and SOC dynamics. These challenges were addressed through scalable methods that balance methodological rigor and real-world applicability.

Chapter 1 introduced the motivation and research objectives, emphasizing the need for interdisciplinary approaches to support tactical-level planning in evolving urban mobility systems. It outlined the limitations of existing planning practices and the potential for integrated modeling to enhance ridership, equity, and infrastructure performance.

Chapter 2 developed a customized GNN framework to predict station-to-station ridership. The model, applied to Toronto’s bike-sharing network, treated the system as a directed weighted graph and outperformed benchmark methods (e.g., XGBoost, ANN) in both predictive accuracy and robustness. Scenario-based expansion planning applications demonstrated the model’s flexibility in supporting transit integration, equity, and network design decisions.

Chapter 3 introduced a force-based CA model to optimize station placement in continuous space. The algorithm accounted for attraction forces from high-demand areas and repulsion or attraction forces between stations to ensure effective spatial distribution. The model achieved near-optimal configurations with reduced computational costs and was validated through a real-world case study in Vancouver.

Chapter 4 addressed electrification planning using a Markovian framework that models SOC transitions and charging dynamics. A two-dimensional birth–death process was formulated to estimate ridership under various charger placement scenarios. A local-improvement heuristic and a single-pooling approximation enabled scalability. Case studies in Pittsburgh, Vancouver, and San Francisco demonstrated the model’s adaptability and operational value in diverse urban contexts.

Chapter 5 translated the research findings into a GIS-based web tool that integrates prediction, optimization, and scenario simulation. The tool was piloted in multiple North American cities and supports strategic planning through interactive visualizations and descriptive, predictive and prescriptive analytics. Its capabilities were demonstrated through an in-depth case study of Toronto, showcasing how the platform supports data-driven scenario planning aligned with local policy goals.

Chapter 6 concludes by summarizing the research contributions and outlining future directions, including demand modeling, multimodal integration, and incorporating behavioral responses and equity objectives.

These chapters offer a comprehensive framework that enhances both the theory and practice of designing bike-sharing systems. In the following section, we summarize the main contributions of this research across modeling, methodology, and implementation.

6.2 Key Contributions

This dissertation responds to growing demands for scalable, data-informed, and electrification-aware planning in station-based bike-sharing systems. It bridges methodological gaps in ridership prediction, spatial optimization, and infrastructure planning, while providing practical tools for decision-makers. The following five contributions summarize the work’s theoretical depth, methodological innovations, and applied relevance.

1. Predicting Origin–Destination Flows with Network-Aware Learning

This research advances beyond traditional station-level inflow and outflow models by capturing the directional structure of ridership flows across an entire network. This approach maintains spatial context while accommodating non-linear interactions between stations by representing the system as a directed, weighted graph and using GraphSAGE-based embeddings. This network-sensitive framework enhances the accuracy of demand prediction. It provides deeper insights into how expansion or modifications to the system, such as adding or removing a station, impact flow patterns throughout the network. It serves as an essential foundation for predictive planning tools that adapt to changes in infrastructure and user behavior.

2. Optimizing Station Locations for Ridership, Not Just Coverage

This work redefines the traditional paradigm of facility location by approaching station placement as a continuous-space optimization problem influenced by two opposing forces: attraction and repulsion. This attraction-repulsion framework creates a more behaviorally informed model of how the arrangement of infrastructure impacts usage. Instead of viewing ridership merely as a result of accessibility or coverage, this approach places it at the center of the optimization process. This redefinition allows cities to design station networks that are both spatially efficient and better aligned with actual travel patterns.

3. Planning Charging Infrastructure through SOC-Aware Optimization

The Markovian modeling framework developed in this research addresses a critical but often overlooked operational constraint in shared e-bike systems: the interaction between battery depletion, charger availability, and ridership. By modeling SOC transitions as part of the network's steady-state behavior, this approach provides a systematic and manageable way to assess how infrastructure decisions impact ridership outcomes. Notably, including approximation strategies allows these insights to be applied to small or theoretical networks and large-scale systems operating under real-world conditions.

4. Delivering a Practical Tool for Strategic Planning

The thesis extends beyond analytical contributions by integrating its models into a decision-support tool for municipal planners. This approach helps bridge the ongoing gap between academic modeling and practical planning. The tool features interactivity, spatial integration, and the ability to simulate future scenarios, making it a valuable platform for adaptive mobility planning. It allows planners to explore trade-offs, such as those between coverage and equity or ridership and cost, enabling them to align infrastructure investments with broader goals of sustainability and accessibility.

5. Toward an Integrated, Scalable Framework for Modern Bike-Share Systems

The final contribution of this research is the integration of three key components: OD ridership prediction, station optimization, and electrification planning, into a unified framework that supports both academic insights and real-world decision-making. Instead of addressing each challenge in isolation, this study connects them within a modular system that reflects the interdependent nature of bike-sharing networks. The models have been designed to interface logically and computationally, allowing for their combined application through a deployable planning tool. This integration ensures that methodological advancements lead to operational benefits, providing cities with a holistic and flexible approach to designing high-performing, future-ready bike-sharing systems.

6.3 Policy and Planning Implications

This research enhances the understanding of bike-sharing systems and how cities design and manage shared mobility infrastructure. First, it redefines what makes a bike-sharing network effective. Instead of relying on heuristic-based expansion or simply responding to demand after it occurs, this work enables cities to plan for latent demand. It helps them anticipate the system-wide effects of local interventions and prioritize investments based on the dynamics of

the network rather than just geographical factors.

Second, the methodology offers a decision-making structure that aligns with broader urban goals, such as decarbonization, modal shifts, and access equity. It allows planners to quantify trade-offs, simulate alternatives, and make transparent decisions amid competing objectives.

Third, the research supports a shift toward adaptive planning practices. The proposed framework supports planning as an ongoing, scenario-based process by integrating optimization and predictive modeling into interactive tools. This approach shifts the focus from static five-year plans to iterative, data-responsive strategies. Finally, the framework addresses a growing need among public agencies: to transform the increasing volume of mobility data into actionable insights. By connecting system analytics with spatial policy design, this work enables cities to turn technical capabilities into strategic foresight.

6.4 Limitations

While this research advances the modeling and planning of station-based bike-sharing systems in several important directions, it is also subject to certain limitations. These relate to modeling assumptions, data availability, and balancing tractability and realism.

First, the ridership prediction model was developed to inform tactical-level decisions, particularly around peak-period demand and network expansion. As such, it intentionally excludes temporal variation across the full operating day and focuses primarily on high-volume OD flows. Low-flow patterns, often relevant for equity assessments, are not emphasized, and behavioral factors such as pricing sensitivity or user preferences are not explicitly modeled. These simplifications improve model focus and interpretability, limiting their applicability for short-term prediction or real-time operational adjustments.

Second, the station optimization model assumes that inflows and outflows at each station are balanced in the long run. While appropriate for strategic planning, this assumption overlooks short-term imbalances that can result in unmet demand or dock saturation. Additionally, the model treats fleet availability as uniform, abstracting from dynamic behaviors such as user routing, rebalancing operations, or temporary capacity constraints. These simplifications help maintain scalability but reduce the model's fidelity in capturing the variability inherent in daily

operations.

Third, the electrification framework models SOC transitions using a Markovian birth–death model, offering a tractable method to evaluate charger placement strategies. However, the model operates under the assumption of uniform charging behavior and independent charging stations. This simplification may overlook important spatial interactions and user decision-making processes. It also fails to account for user preferences related to SOC levels, which can influence bike availability and trip completion rates. Although these assumptions were necessary for computational feasibility, they may limit the framework’s applicability in highly dynamic or behaviorally heterogeneous settings.

Finally, while the models were validated across various cities of different sizes and typologies, the framework assumes a centralized, station-based system. It may require adaptation for hybrid or dockless systems. Transferability to other geographies also depends on the availability and quality of input data, including OD flows, station metadata, and contextual variables such as demographics and land use.

While these limitations do not diminish the research’s contributions, they highlight important opportunities for refinement and extension, many of which are outlined in the subsequent section on future research directions.

6.5 Directions for Future Research

The work presented in this dissertation lays a strong foundation for data-driven, electrification-aware planning of station-based bike-sharing systems. Building on this foundation, several avenues offer promising directions for future research and methodological refinement.

Incorporating Temporal and Behavioral Dynamics: Current models operate under steady-state or peak-period assumptions to support strategic planning. Extending the framework to incorporate time-varying demand, multi-period scheduling, and user behavior, including responsiveness to pricing, incentives, or station amenities, would increase realism and broaden the framework’s applicability to operational planning and policy evaluation. Temporal models could also support demand prediction across seasons or special events, improving resource allocation.

Enhancing Equity and Accessibility Modeling: While the framework supports scenario analysis for equitable access, future work could embed explicit equity objectives within the optimization process. This includes integrating accessibility metrics, service equity constraints, or demographic-weighted ridership goals. Linking predictive models to measures of transport justice would better align system design with the priorities of inclusive urban mobility.

Integration with Multimodal and Intermodal Systems: Future research should explore extensions of the framework to multimodal networks that include transit, walking, and other micromobility modes. This would enable co-optimization of bike-sharing stations with transit stops or mobility hubs, supporting integrated first-/last-mile planning. It would also allow for evaluation of policy strategies such as fare integration, cross-modal incentives, or infrastructure co-location.

Extensions to Dockless and Hybrid Systems: The current models are tailored to docked systems with fixed infrastructure. As cities adopt hybrid or fully dockless systems, future research could generalize the spatial framework to accommodate fluid vehicle distributions, virtual stations, and user-defined drop zones.

Long-Term Policy Simulation and System Design: Finally, future work could leverage the framework to simulate long-term system evolution under different policy, investment, or behavioral scenarios. This includes evaluating trade-offs within budget constraints, testing scenarios for carbon reduction or health impacts, and developing phased implementation strategies. By combining the planning models with simulations or system dynamics models, we can gain a more comprehensive understanding of the role of bike-sharing in future urban transportation ecosystems.

These directions collectively suggest a future for shared mobility planning that is more dynamic, responsive, and integrated, especially for cycling-based systems. By expanding the current framework to incorporate aspects such as temporal complexity, multimodal coordination, and social equity, future research can enhance the role of bike-sharing systems as a vital element of sustainable urban transportation. The continued evolution of these tools and models will be essential to support cities in navigating the challenges of electrification, climate adaptation, and inclusive mobility design.

6.6 Final Remarks

Shared micro-mobility systems, especially bike-sharing networks, present a significant opportunity to rethink urban transportation in ways that are flexible, inclusive, and sustainable. However, to fully realize this potential, we need more than just expanding infrastructure; we must also gain a deeper understanding of how demand, space, and operations interact, as well as how technology can be leveraged to make smarter decisions.

This dissertation contributes to that vision by proposing a comprehensive framework for planning station-based bike-sharing systems. This framework integrates ridership prediction, spatial optimization, and strategies that consider electrification. Through methodological innovations and practical models, this work provides a scalable approach to designing networks that are high-performing and future-ready. The developed models enhance academic knowledge and offer practical insights for planners and policymakers aiming to make urban mobility more resilient, equitable, and efficient.

As data, machine learning, and optimization increasingly converge, new opportunities arise for transportation planning. Decisions can now be simulated, tested, and refined on an ongoing basis. The tools and concepts created in this research are designed to adapt to this evolving landscape, facilitating the shift from static planning to intelligent, adaptive systems that benefit both people and the planet.

Bibliography

(2020). Mapbox api. [Online]. Available: <https://docs.mapbox.com/api/overview/>.

(2023). Citi bike monthly operating reports. [Online]. Available: <https://citibikenyc.com/system-data/operating-reports>.

Akturk, D., Candogan, O., and Gupta, V. (2024). Managing resources for shared micromobility: Approximate optimality in large-scale systems. *Management Science*.

Ansari, S., Bacsdere, M., Li, X., Ouyang, Y., and Smilowitz, K. (2018). Advancements in continuous approximation models for logistics and transportation systems: 1996–2016. *Transportation Research Part B: Methodological*, 107:229–252.

Bahadori, M. S., Goncalves, A. B., and Moura, F. (2021). A systematic review of station location techniques for bicycle-sharing systems planning and operation. *ISPRS International Journal of Geo-Information*, 10(8):554.

Banerjee, S., Freund, D., and Lykouris, T. (2022). Pricing and optimization in shared vehicle systems: An approximation framework. *Operations Research*, 70(3):1783–1805.

Banerjee, S., Kabir, M. M., Khadem, N. K., and Chavis, C. (2020). Optimal locations for bike-share stations: A new gis based spatial approach. *Transportation Research Interdisciplinary Perspectives*, 4:100101.

Barth, M. and Todd, M. (1999). Simulation model performance analysis of a multiple station shared vehicle system. *Transportation Research Part C: Emerging Technologies*, 7(4):237–259.

- 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.
- Buck, D. and Buehler, R. (2012). Bike lanes and other determinants of capital bikeshare trips. In *91st Transportation research board annual meeting*, pages 703–706.
- Caggiani, L., Colovic, A., and Ottomanelli, M. (2020). An equality-based model for bike-sharing stations location in bicycle-public transport multimodal mobility. *Transportation Research Part A: Policy and Practice*, 140:251–265.
- Caicedo, A., Estrada, M., Medina-Tapia, M., and Mayorga, M. (2023). Optimizing bike network design: A cost-effective methodology for heterogeneous travel demands using continuous approximation techniques. *Transportation research part A: policy and practice*, 176:103826.
- Chai, D., Wang, L., and Yang, Q. (2018). Bike flow prediction with multi-graph convolutional networks. In *Proceedings of the 26th ACM SIGSPATIAL international conference on advances in geographic information systems*, pages 397–400.
- Chen, Y. and Liu, Y. (2023). Integrated optimization of planning and operations for shared autonomous electric vehicle systems. *Transportation Science*, 57(1):106–134.
- 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. *Applied Soft Computing*, 67:519–528.
- Chen, Z., Hu, Y., Li, J., and Wu, X. (2020). Optimal deployment of electric bicycle sharing stations: model formulation and solution technique. *Networks and Spatial Economics*, 20:99–136.
- Chiou, Y.-C. and Wu, K.-C. (2024). Bikesharing: The first-and last-mile service of public transportation? evidence from an origin–destination perspective. *Transportation Research Part A: Policy and Practice*, 187:104162.
- Daganzo, C. F. (1984). The length of tours in zones of different shapes. *Transportation Research Part B: Methodological*, 18(2):135–145.

- Daganzo, C. F. (2010). Structure of competitive transit networks. *Transportation Research Part B: Methodological*, 44(4):434–446.
- Daganzo, C. F. and Newell, G. F. (1986). Configuration of physical distribution networks. *Networks*, 16(2):113–132.
- Datner, S., Raviv, T., Tzur, M., and Chemla, D. (2019). Setting inventory levels in a bike sharing network. *Transportation Science*, 53(1):62–76.
- Dey, B. K., Anowar, S., and Eluru, N. (2021). A framework for estimating bikeshare origin destination flows using a multiple discrete continuous system. *Transportation Research Part A: Policy and Practice*, 144:119–133.
- Ding, H., Lu, Y., Sze, N., and Li, H. (2022). Effect of dockless bike-sharing scheme on the demand for london cycle hire at the disaggregate level using a deep learning approach. *Transportation research part A: policy and practice*, 166:150–163.
- Drezner, Z. and Hamacher, H. W. (2004). *Facility location: applications and theory*. Springer Science & Business Media.
- Fournier, N., Christofa, E., and Knodler Jr, M. A. (2017). A sinusoidal model for seasonal bicycle demand estimation. *Transportation research part D: transport and environment*, 50:154–169.
- Frade, I. and Ribeiro, A. (2014). Bicycle sharing systems demand. *Procedia-Social and Behavioral Sciences*, 111:518–527.
- Frade, I. and Ribeiro, A. (2015). Bike-sharing stations: A maximal covering location approach. *Transportation Research Part A: Policy and Practice*, 82:216–227.
- Freund, D., Henderson, S. G., and Shmoys, D. B. (2017). Minimizing multimodular functions and allocating capacity in bike-sharing systems. In *Integer Programming and Combinatorial Optimization: 19th International Conference, IPCO 2017, Waterloo, ON, Canada, June 26-28, 2017, Proceedings 19*, pages 186–198. Springer.

- Fricker, C. and Gast, N. (2016). Incentives and redistribution in homogeneous bike-sharing systems with stations of finite capacity. *Euro journal on transportation and logistics*, 5:261–291.
- Froehlich, J., Neumann, J., Oliver, N., et al. (2009). Sensing and predicting the pulse of the city through shared bicycling. In *IJCAI*, volume 9, pages 1420–1426.
- Guidon, S., Reck, D. J., and Axhausen, K. (2020). Expanding a (n)(electric) bicycle-sharing system to a new city: Prediction of demand with spatial regression and random forests. *Journal of Transport Geography*, 84:102692.
- Guo, R., Jiang, Z., Huang, J., Tao, J., Wang, C., Li, J., and Chen, L. (2019). Bikenet: Accurate bike demand prediction using graph neural networks for station rebalancing. In *2019 IEEE SmartWorld, Ubiquitous Intelligence & Computing, Advanced & Trusted Computing, Scalable Computing & Communications, Cloud & Big Data Computing, Internet of People and Smart City Innovation (SmartWorld/SCALCOM/UIC/ATC/CBDCom/IOP/SCI)*, pages 686–693. IEEE.
- Ham, S. W., Cho, J.-H., Park, S., and Kim, D.-K. (2021). Spatiotemporal demand prediction model for e-scooter sharing services with latent feature and deep learning. *Transportation research record*, 2675(11):34–43.
- Hamilton, W., Ying, Z., and Leskovec, J. (2017). Inductive representation learning on large graphs. *Advances in neural information processing systems*, 30.
- He, L., Ma, G., Qi, W., and Wang, X. (2021a). Charging an electric vehicle-sharing fleet. *Manufacturing & Service Operations Management*, 23(2):471–487.
- He, P., Zheng, F., Belavina, E., and Girotra, K. (2021b). Customer preference and station network in the london bike-share system. *Management Science*, 67(3):1392–1412.
- Hosseinzadeh, A., Karimpour, A., and Kluger, R. (2021). Factors influencing shared micromobility services: An analysis of e-scooters and bikeshare. *Transportation Research Part D: Transport and Environment*, 100:103047.

- Hyland, M., Hong, Z., de Farias Pinto, H. K. R., and Chen, Y. (2018). Hybrid cluster-regression approach to model bikeshare station usage. *Transportation Research Part A: Policy and Practice*, 115:71–89.
- Jiang, W. (2022). Bike sharing usage prediction with deep learning: a survey. *Neural Computing and Applications*, 34(18):15369–15385.
- Kabra, A., Belavina, E., and Girotra, K. (2020). Bike-share systems: Accessibility and availability. *Management Science*, 66(9):3803–3824.
- Kim, S., Choo, S., Lee, G., and Kim, S. (2022). Predicting demand for shared e-scooter using community structure and deep learning method. *Sustainability*, 14(5):2564.
- Kim, T. S., Lee, W. K., and Sohn, S. Y. (2019). Graph convolutional network approach applied to predict hourly bike-sharing demands considering spatial, temporal, and global effects. *PloS one*, 14(9):e0220782.
- Kou, Z. and Cai, H. (2021). Incorporating spatial network information to improve demand prediction for bike share system expansion. In *Proceedings of the 10th International Workshop on Urban Computing. ACM, Beijing*.
- Kumar, S., Spezzano, F., Subrahmanian, V., and Faloutsos, C. (2016). Edge weight prediction in weighted signed networks. In *2016 IEEE 16th International Conference on Data Mining (ICDM)*, pages 221–230. IEEE.
- Lai, Z. and Li, S. (2024). Towards a multimodal charging network: Joint planning of charging stations and battery swapping stations for electrified ride-hailing fleets. *Transportation Research Part B: Methodological*, 183:102928.
- Larsen, J., Patterson, Z., and El-Geneidy, A. (2013). Build it. but where? the use of geographic information systems in identifying locations for new cycling infrastructure. *International Journal of Sustainable Transportation*, 7(4):299–317.
- Li, G., Wang, X., Njoo, G. S., Zhong, S., Chan, S.-H. G., Hung, C.-C., and Peng, W.-C. (2022).

- A data-driven spatial-temporal graph neural network for docked bike prediction. In *2022 IEEE 38th International Conference on Data Engineering (ICDE)*, pages 713–726. IEEE.
- Li, Q., Luca, D., Fuerst, F., and Wei, Z. (2024). Success in tandem? the impact of the introduction of e-bike sharing on bike sharing usage. *Research in Transportation Economics*, 107:101476.
- Li, Q.-L., Chen, C., Fan, R.-N., Xu, L., and Ma, J.-Y. (2016a). Queueing analysis of a large-scale bike sharing system through mean-field theory. *arXiv preprint arXiv:1603.09560*.
- Li, X., Ma, J., Cui, J., Ghiasi, A., and Zhou, F. (2016b). Design framework of large-scale one-way electric vehicle sharing systems: A continuum approximation model. *Transportation Research Part B: Methodological*, 88:21–45.
- Liang, Y., Ding, F., Huang, G., and Zhao, Z. (2023). Deep trip generation with graph neural networks for bike sharing system expansion. *arXiv preprint arXiv:2303.11977*.
- Lin, L., He, Z., and Peeta, S. (2018). Predicting station-level hourly demand in a large-scale bike-sharing network: A graph convolutional neural network approach. *Transportation Research Part C: Emerging Technologies*, 97:258–276.
- Liu, J., Chen, W., and Sun, L. (2024). A data-driven optimization framework for static rebalancing operations in bike sharing systems. *INFORMS Journal on Computing*.
- Liu, J., Li, Q., Qu, M., Chen, W., Yang, J., Xiong, H., Zhong, H., and Fu, Y. (2015). Station site optimization in bike sharing systems. In *2015 IEEE International Conference on Data Mining*, pages 883–888. IEEE.
- Liu, J., Sun, L., Li, Q., Ming, J., Liu, Y., and Xiong, H. (2017a). Functional zone based hierarchical demand prediction for bike system expansion. In *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 957–966.
- Liu, P., Ding, Y., Fu, T., and Li, D. (2017b). Joint resource optimization of shared electric bicycles in smart cities. In *2017 IEEE International Symposium on Parallel and Distributed*

Processing with Applications and 2017 IEEE International Conference on Ubiquitous Computing and Communications (ISPA/IUCC), pages 1138–1143. IEEE.

Liu, X., Yu, J., Zhao, J., and Schneider, R. J. (2025). Bikeshare impacts on bus ridership: Unraveling the rail proximity effect. *Transportation Research Part A: Policy and Practice*, 192:104373.

Liu, Y. and Ouyang, Y. (2025). Planning charging stations for mixed docked and dockless operations of shared electric micromobility systems. *Transportation Research Part C: Emerging Technologies*, 171:104989.

Luo, X., Gu, W., and Fan, W. (2021). Joint design of shared-bike and transit services in corridors. *Transportation Research Part C: Emerging Technologies*, 132:103366.

Ma, X., Yin, Y., Jin, Y., He, M., and Zhu, M. (2022). Short-term prediction of bike-sharing demand using multi-source data: a spatial-temporal graph attentional lstm approach. *Applied Sciences*, 12(3):1161.

Martinez, L. M., Caetano, L., Eiró, T., and Cruz, F. (2012). An optimisation algorithm to establish the location of stations of a mixed fleet biking system: an application to the city of lisbon. *Procedia-Social and Behavioral Sciences*, 54:513–524.

Meddin, R., DeMaio, P., O’Brien, O., Rabello, R., Yu, C., Seamon, J., and Benicchio, T. (2022). The meddin bike-sharing world map. [Online]. Available: <https://bikesharingworldmap.com/>.

Müller, F. (2023). Link and edge weight prediction in air transport networks—an rnn approach. *Physica A: Statistical Mechanics and its Applications*, page 128490.

Nair, R. and Miller-Hooks, E. (2011). Fleet management for vehicle sharing operations. *Transportation Science*, 45(4):524–540.

National Association of City Transportation Officials (2024). Shared micromobility in 2023.

- Neumann-Saavedra, B. A., Crainic, T. G., Gendron, B., Mattfeld, D. C., and Römer, M. (2020). Integrating resource management in service network design for bike-sharing systems. *Transportation Science*, 54(5):1251–1271.
- Newell, G. F. (1971). Dispatching policies for a transportation route. *Transportation Science*, 5(1):91–105.
- Newell, G. F. (1973). Scheduling, location, transportation, and continuum mechanics: some simple approximations to optimization problems. *SIAM Journal on Applied Mathematics*, 25(3):346–360.
- Newman, M. (2018). *Networks*. Oxford university press.
- Nikiforiadis, A., Aifadopoulou, G., Grau, J. M. S., and Boufidis, N. (2021). Determining the optimal locations for bike-sharing stations: methodological approach and application in the city of thessaloniki, greece. *Transportation Research Procedia*, 52:557–564.
- Orvin, M. M., Bachhal, J. K., and Fatmi, M. R. (2022). Modeling the demand for shared e-scooter services. *Transportation research record*, 2676(3):429–442.
- Ouyang, Y. (2007). Design of vehicle routing zones for large-scale distribution systems. *Transportation Research Part B: Methodological*, 41(10):1079–1093.
- Ouyang, Y. and Daganzo, C. F. (2006). Discretization and validation of the continuum approximation scheme for terminal system design. *Transportation Science*, 40(1):89–98.
- Ouyang, Y., Nourbakhsh, S. M., and Cassidy, M. J. (2014). Continuum approximation approach to bus network design under spatially heterogeneous demand. *Transportation Research Part B: Methodological*, 68:333–344.
- Philips, I., Anable, J., and Chatterton, T. (2022). E-bikes and their capability to reduce car co2 emissions. *Transport Policy*, 116:11–23.
- Qian, X., Jaller, M., and Circella, G. (2022). Equitable distribution of bikeshare stations: An optimization approach. *Journal of transport geography*, 98:103174.

- Rahmani, S., Baghbani, A., Bouguila, N., and Patterson, Z. (2023). Graph neural networks for intelligent transportation systems: A survey. *IEEE Transactions on Intelligent Transportation Systems*.
- Raviv, T. and Kolka, O. (2013). Optimal inventory management of a bike-sharing station. *Iie Transactions*, 45(10):1077–1093.
- Regue, R. and Recker, W. (2014). Proactive vehicle routing with inferred demand to solve the bikesharing rebalancing problem. *Transportation Research Part E: Logistics and Transportation Review*, 72:192–209.
- Rixey, R. A. (2013). Station-level forecasting of bikesharing ridership: Station network effects in three us systems. *Transportation research record*, 2387(1):46–55.
- Roni, M. S., Yi, Z., and Smart, J. G. (2019). Optimal charging management and infrastructure planning for free-floating shared electric vehicles. *Transportation Research Part D: Transport and Environment*, 76:155–175.
- Samet, B., Couffin, F., Zolghadri, M., Barkallah, M., and Haddar, M. (2018). Model reduction for studying a bike sharing system as a closed queuing network. *Procedia Manufacturing*, 25:39–46.
- Saum, N., Sugiura, S., and Piantanakulchai, M. (2020). Short-term demand and volatility prediction of shared micro-mobility: a case study of e-scooter in thammasat university. In *2020 Forum on Integrated and Sustainable Transportation Systems (FISTS)*, pages 27–32. IEEE.
- Sohrabi, S., Paleti, R., Balan, L., and Cetin, M. (2020). Real-time prediction of public bike sharing system demand using generalized extreme value count model. *Transportation Research Part A: Policy and Practice*, 133:325–336.
- Song, J., Li, B., Szeto, W., and Zhan, X. (2024). A station location design problem in a bike-sharing system with both conventional and electric shared bikes considering bike users’ roaming delay costs. *Transportation research part E: logistics and transportation review*, 181:103350.

- Song, J.-C., Hsieh, I.-Y. L., and Chen, C.-S. (2023). Sparse trip demand prediction for shared e-scooter using spatio-temporal graph neural networks. *Transportation research part D: transport and environment*, 125:103962.
- Soriguera, F. and Jiménez-Meroño, E. (2020). A continuous approximation model for the optimal design of public bike-sharing systems. *Sustainable Cities and Society*, 52:101826.
- Tavares, B., Soares, C., and Marques, M. (2022). A cluster-based trip prediction graph neural network model for bike sharing systems. *arXiv preprint arXiv:2201.00720*.
- Thomas, D., Klonari, V., Vallée, F., and Ioakimidis, C. S. (2015). Implementation of an e-bike sharing system: The effect on low voltage network using pv and smart charging stations. In *2015 International Conference on Renewable Energy Research and Applications (ICRERA)*, pages 572–577. IEEE.
- Thomas, T., Jaarsma, C., and Tutert, S. (2009). Temporal variations of bicycle demand in the netherlands: The influence of weather on cycling.
- Wang, J., Tsai, C.-H., and Lin, P.-C. (2016). Applying spatial-temporal analysis and retail location theory to public bikes site selection in taipei. *Transportation Research Part A: Policy and Practice*, 94:45–61.
- Wang, S., Xie, H., Yun, B., Pu, X., and Qiu, Z. (2024). Optimization strategy for the spatiotemporal layout of e-bike charging piles from the perspective of sustainable campus planning: A case study of zijingang campus of zhejiang university. *Sustainability*, 16(13):5690.
- Wang, Y. and Szeto, W. (2021). The dynamic bike repositioning problem with battery electric vehicles and multiple charging technologies. *Transportation Research Part C: Emerging Technologies*, 131:103327.
- Xiao, G., Wang, R., Zhang, C., and Ni, A. (2021). Demand prediction for a public bike sharing program based on spatio-temporal graph convolutional networks. *Multimedia Tools and Applications*, 80:22907–22925.

- Xiao, Y., Ahmed, F., and Sha, Z. (2022). Travel links prediction in shared mobility networks using graph neural network models. In *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, volume 86212, page V002T02A079. American Society of Mechanical Engineers.
- Xiao, Y., Ahmed, F., and Sha, Z. (2023). Graph neural network-based design decision support for shared mobility systems. *Journal of Mechanical Design*, pages 1–17.
- Xie, X., Dai, X., and Pei, Z. (2024a). Empowering the capillary of the urban daily commute: Battery deployment analysis for the locker-based e-bike battery swapping. *Transportation Science*, 58(1):176–197.
- Xie, Y., Smart, M., and Noland, R. B. (2024b). Powering bikeshare in new york city: does the usage of e-bikes differ from regular bikes? *Transportation Planning and Technology*, pages 1–28.
- Xu, C., Ji, J., and Liu, P. (2018). The station-free sharing bike demand forecasting with a deep learning approach and large-scale datasets. *Transportation research part C: emerging technologies*, 95:47–60.
- Xu, M., Di, Y., Yang, H., Chen, X., and Zhu, Z. (2023). Multi-task supply-demand prediction and reliability analysis for docked bike-sharing systems via transformer-encoder-based neural processes. *Transportation Research Part C: Emerging Technologies*, 147:104015.
- Yang, H., Qin, W., and Zhou, Y. (2022). User-based battery swapping strategy in an electric bike-sharing system. In *2022 IEEE International Conference on Industrial Engineering and Engineering Management (IEEM)*, pages 0236–0240. IEEE.
- Yang, Y., Heppenstall, A., Turner, A., and Comber, A. (2020). Using graph structural information about flows to enhance short-term demand prediction in bike-sharing systems. *Computers, Environment and Urban Systems*, 83:101521.
- Yang, Z., Chen, J., Hu, J., Shu, Y., and Cheng, P. (2019). Mobility modeling and data-driven closed-loop prediction in bike-sharing systems. *IEEE Transactions on Intelligent Transportation Systems*, 20(12):4488–4499.

- Zhan, Z., Guo, Y., Noland, R. B., He, S. Y., and Wang, Y. (2023). Analysis of links between dockless bikeshare and metro trips in beijing. *Transportation Research Part A: Policy and Practice*, 175:103784.
- Zhang, J., Meng, M., Wang, D. Z., Zhou, L., and Han, L. (2023). Optimal bike allocations in a competitive bike sharing market. *Journal of Cleaner Production*, 384:135602.
- Zhao, L., Shen, S., and Zhao, Z. (2024). Planning decentralized battery-swapping recharging facilities for e-bike sharing systems. *Sustainable Cities and Society*, 101:105118.
- Zhou, J., Cui, G., Hu, S., Zhang, Z., Yang, C., Liu, Z., Wang, L., Li, C., and Sun, M. (2020). Graph neural networks: A review of methods and applications. *AI open*, 1:57–81.
- Zhou, Y., Lin, Z., Guan, R., and Sheu, J.-B. (2023). Dynamic battery swapping and rebalancing strategies for e-bike sharing systems. *Transportation Research Part B: Methodological*, 177:102820.
- Zhu, S. (2021). Optimal fleet deployment strategy: Model the effect of shared e-bikes on bike-sharing system. *Journal of Advanced Transportation*, 2021(1):6678637.
- Zhu, Z., Xu, M., Di, Y., and Yang, H. (2022). Fitting spatial-temporal data via a physics regularized multi-output grid gaussian process: case studies of a bike-sharing system. *IEEE Transactions on Intelligent Transportation Systems*, 23(11):21090–21101.

Appendices

A Appendix: List of Notations for Chapter 3

Symbol	Description
$X = \{x_1, x_2, \dots, x_N\}$	Set of bike-sharing station locations
S	Continuous service area
$S' \subseteq S$	Inner region ensuring full containment of influence areas
q_i	Total dock capacity of station i
b_i	Number of available bikes at station i
A_i	Influence area of station i
$ A_i $	Area of the influence region of station i
$A_{\max}$	Maximum allowable influence area
r_i	Radius of the influence disk for station i
$r_{\max}$	Maximum allowable radius of influence
$\Lambda_I(x)$	Origin (trip production) density at location x
$\Lambda_J(x)$	Destination (trip attraction) density at location x
$\Lambda_{I,\text{avg},i}$	Average origin density in station i 's influence area
$\Lambda_{J,\text{avg},i}$	Average destination density in station i 's influence area
$P_i(X)$	Total trip production associated with station i
$T_i(X)$	Total trip termination associated with station i
$P_i^{\text{initial}}, T_i^{\text{initial}}$	Initial trip production and termination estimates
λ_{ij}	Flow (ridership) between stations i and j
$C_{ij}(x_i, x_j)$	Impedance function between stations i and j
d_{ij}	Distance between stations i and j
D_{opt}	Optimal inter-station distance for maximizing ridership
$D_{\min}$	Minimum allowable inter-station distance
$F_{SS_{ij}}$	Station-to-station force (repulsion or attraction)

Symbol	Description
F_{ACij}	Accessibility force (repulsion)
F_{Oi}, F_{Di}	Demand-to-station forces from origins/destinations (attraction)
η	Balance between inter-station and demand-to-station forces
μ	Step size for station movement
Δ_i	Movement magnitude of station i in one iteration
Δ_{avg}	Average movement across all stations
ϵ	Convergence threshold
α	Distance decay exponent in attraction force
β	Capacity sensitivity exponent in attraction force
K	Strength coefficient in the accessibility force

B Appendix: Global balance equations and steady-state solutions for Small C_s

This appendix provides the global balance equations and steady-state probabilities for small C_s to illustrate how system complexity scales with station capacity. Unlike one-dimensional birth-death processes, which often yield product-form solutions, our two-dimensional Markov process introduces intricate state dependencies due to arrivals, departures, and charging events. As a result, explicit analytical solutions grow complex, and no general closed-form expression in terms of C_s emerges. With the state space expanding quadratically based on Equation 4.5, analytical tractability diminishes rapidly.

Case $C_s = 1$ In the simplest case with $C_s = 1$, the state space is limited to only 3 states with 5 transitions. Referring to the balance equations in the main text, we find the steady-state probabilities $\pi = [\pi_{00}, \pi_{01}, \pi_{10}]$:

$$\begin{aligned}\pi_{00} &= \frac{\mu_s^H(\gamma_s + \mu_s^L)}{\lambda_s^H(\gamma_s + \mu_s^L) + \lambda_s^L(\gamma_s + \mu_s^H) + \mu_s^H(\gamma_s + \mu_s^L)}, \\ \pi_{01} &= \frac{\gamma_s \lambda_s^L + \lambda_s^H(\gamma_s + \mu_s^L)}{\lambda_s^H(\gamma_s + \mu_s^L) + \lambda_s^L(\gamma_s + \mu_s^H) + \mu_s^H(\gamma_s + \mu_s^L)}, \\ \pi_{10} &= \frac{\lambda_s^L \mu_s^H}{\lambda_s^H(\gamma_s + \mu_s^L) + \lambda_s^L(\gamma_s + \mu_s^H) + \mu_s^H(\gamma_s + \mu_s^L)}.\end{aligned}$$

Case $C_s = 2$ For a station capacity of $C_s = 2$, the state space comprises 6 distinct states. The resulting Markov chain has a total of 15 transitions, which can be thought of as 5 primary transitions complemented by 10 additional ones arising from the interplay of multiple event types (see Figure 4.3-a). This case illustrates how even a modest increase in capacity substantially increases the system's complexity compared to the $C_s = 1$ scenario.

$$\pi_{10} = \frac{\lambda_s^L (2\gamma_s \mu_s^{H2} + 2\mu_s^{H2} \mu_s^L)}{\lambda_s^{H2} (\gamma_s^2 + 2\gamma_s \mu_s^L + \mu_s^{L2}) + \lambda_s^H (2\gamma_s^2 \lambda_s^L + 2\gamma_s^2 \mu_s^H + 2\gamma_s \lambda_s^L \mu_s^H + 2\gamma_s \lambda_s^L \mu_s^L + 4\gamma_s \mu_s^H \mu_s^L + 2\lambda_s^L \mu_s^H \mu_s^L + 2\mu_s^H \mu_s^{L2}) + \lambda_s^{L2} (\gamma_s^2 + 2\gamma_s \mu_s^H + \mu_s^{H2}) + \lambda_s^L (2\gamma_s^2 \mu_s^H + 2\gamma_s \mu_s^{H2} + 2\gamma_s \mu_s^H \mu_s^L + 2\mu_s^{H2} \mu_s^L) + \mu_s^{H2} (2\gamma_s^2 + 4\gamma_s \mu_s^L + 2\mu_s^{L2})},$$

Case $C_s = 3$ For $C_s = 3$, the state space consists of ten states of $\{(0,0), (0,1), (0,2), (0,3), (1,0), (1,1), (1,2), (2,0), (2,1), (3,0)\}$, as illustrated in Figure 4.1. As the station capacity increases, the state space expands quadratically, and the corresponding steady-state probabilities become exceedingly complex. For example, the steady-state probability for the state $(0,0)$ is given by:

$$\pi_{00} = \frac{\mu_s^{H3} (6\gamma_s^3 + 18\gamma_s^2 \mu_s^L + 18\gamma_s \mu_s^{L2} + 6\mu_s^{L3})}{D_{00}}$$

$$\begin{aligned} D_{00} = & \lambda_s^{H3} (\gamma_s^3 + 3\gamma_s^2 \mu_s^L + 3\gamma_s (\mu_s^L)^2 + (\mu_s^L)^3) + \lambda_s^{H2} (3\gamma_s^3 \lambda_s^L + 3\gamma_s^3 \mu_s^H + 3\gamma_s^2 \lambda_s^L \mu_s^H + 6\gamma_s^2 \lambda_s^L \mu_s^L + 9\gamma_s^2 \mu_s^H \mu_s^L + 6\gamma_s \lambda_s^L \mu_s^H \mu_s^L + 3\gamma_s \lambda_s^L \mu_s^{L2} + 9\gamma_s \mu_s^H \mu_s^{L2} + 6\gamma_s \lambda_s^L \mu_s^H \mu_s^L + 3\gamma_s \lambda_s^L \mu_s^{L2} + 9\gamma_s \mu_s^H \mu_s^{L2} + 3\lambda_s^L \mu_s^H \mu_s^{L2} + 3\mu_s^H \mu_s^{L3}) \\ & + \lambda_s^H (3\gamma_s^3 \lambda_s^{L2} + 6\gamma_s^3 \lambda_s^L \mu_s^H + 6\gamma_s^3 \mu_s^{H2} + 6\gamma_s^2 \lambda_s^{L2} \mu_s^H + 3\gamma_s^2 \lambda_s^{L2} \mu_s^L + 6\gamma_s^2 \lambda_s^L \mu_s^{H2} + 12\gamma_s^2 \lambda_s^L \mu_s^H \mu_s^L + 18\gamma_s^2 \mu_s^{H2} \mu_s^L + 3\gamma_s \lambda_s^{L2} \mu_s^{H2} + 6\gamma_s \lambda_s^{L2} \mu_s^H \mu_s^L + 12\gamma_s \lambda_s^L \mu_s^{H2} \mu_s^L + 6\gamma_s \lambda_s^L \mu_s^H \mu_s^{L2} + 18\gamma_s \mu_s^{H2} \mu_s^{L2} + 3\lambda_s^{L2} \mu_s^{H2} \mu_s^L + 6\lambda_s^L \mu_s^{H2} \mu_s^{L2} + 6\mu_s^{H2} \mu_s^{L3}) \\ & + \lambda_s^{L3} (\gamma_s^3 + 3\gamma_s^2 \mu_s^H + 3\gamma_s \mu_s^{H2} + \mu_s^{H3}) + \lambda_s^{L2} (3\gamma_s^3 \mu_s^H + 6\gamma_s^2 \mu_s^{H2} + 3\gamma_s^2 \mu_s^H \mu_s^L + 3\gamma_s \mu_s^{H3} + 6\gamma_s \mu_s^{H2} \mu_s^L + 3\mu_s^{H3} \mu_s^L) + \lambda_s^L (6\gamma_s^3 \mu_s^{H2} + 6\gamma_s^2 \mu_s^{H3} + 12\gamma_s^2 \mu_s^{H2} \mu_s^L + 12\gamma_s \mu_s^{H3} \mu_s^L + 6\gamma_s \mu_s^{H2} \mu_s^{L2} + 6\mu_s^{H3} \mu_s^{L2}) + \mu_s^{H3} (6\gamma_s^3 + 18\gamma_s^2 \mu_s^L + 18\gamma_s \mu_s^{L2} + 6\mu_s^{L3}) \end{aligned}$$

The equations include complex interplays of arrival, departure, and charging parameters, reflecting the intricate dependencies within the system. Various simplification approaches, including factorization and term grouping, have been explored to derive a more compact representation. However, as C_s , no discernible pattern or closed-form expression emerges, highlighting the inherent analytical intractability of the steady-state probabilities in this model.

C Appendix: Shared E-Scooter Systems

This appendix summarizes a separate but thematically related study focused on shared e-scooter systems. While the core of this dissertation addresses station-based bike-sharing, shared e-scooters represent another vital component of the micromobility landscape. These systems are influenced by similar spatial, demographic, and infrastructural factors and contribute to the same broader goals of sustainable and flexible urban mobility. Unlike docked bike-sharing, shared e-scooter systems often operate as hybrid or dockless networks. As such, our approach in this study centered on tactical and operational planning rather than long-term infrastructure placement. Specifically, the focus shifted from total ridership prediction to accurate estimation of pickup and drop-off volumes across space and time, an essential component for managing fleet distribution, ensuring service availability, and minimizing operational disruptions.

C.1 Data and Modeling Framework

The study used data from SCOOTY's shared e-scooter program in Brampton, Ontario. The dataset included:

- **Ridership Data:** Geolocated pickup and drop-off records, enriched with ride-level attributes such as time, distance, and duration.
- **Geofencing Data:** Spatial boundaries for restricted, parking, slow-speed, and service areas.
- **Transit Data:** Boarding and alighting volumes at bus stops, supporting analysis of inter-modal integration.
- **Census Data:** Demographic indicators (e.g., population, income, age)
- **POIs from API:** Land use data (e.g., commercial centers, community hubs).
- **Weather APIs:** Historical temperature and precipitation data for temporal enrichment of the demand models.

Consistent with the methodology applied in earlier chapters, this study used a hexagonal grid spatial framework, with each hex cell representing a localized unit of analysis. Pickup and drop-off events and all background features were aggregated at the cell level to ensure uniformity across modeling components. Previous studies have explored shared e-scooter demand through diverse modeling approaches. Deep learning methods have become more popular, including context-aware LSTM models, spatiotemporal recurrent frameworks, and graph neural networks. These approaches take into account factors such as POIs, weather conditions, land use, and temporal activity patterns (Kim et al., 2022; Ham et al., 2021; Saum et al., 2020; Song et al., 2023). These models aim to accurately capture the dynamic and spatially varied nature of micromobility usage. Interpretable statistical methods have also been used. For example, Orvin et al. (2022) used a binomial model to tackle issues related to sparsity and overdispersion in trip data. This model also helped identify significant predictors, such as the built environment and time-of-day effects.

Building on these insights, we developed two predictive models: a short-term daily model and a longer-horizon seasonal model. Both models operate at the node level, using a hexagonal grid to predict localized pickup and drop-off volumes rather than OD flows. A systematic training pipeline was used, including cross-validation and hyperparameter tuning via grid search. Several modeling approaches were benchmarked; XGBoost consistently outperformed alternatives, yielding reliable RMSE values across planning horizons. As more granular and temporally rich data becomes available, the predictive accuracy of these models is expected to improve further, enhancing their utility in dynamic micromobility planning. The models are summarized as follows:

- 1) Daily Prediction Model:** Designed for short-term operational responsiveness, this model incorporates temporal features (e.g., day type, hour), weather variables, spatial and demographic features (e.g., population density, income, POIs, transit proximity), , travel behavior (e.g., recent trip patterns), and historical demand activity, including prior pickups/drop-offs within each cell.
- 2) Seasonal Prediction Model:** Tailored for longer-term planning horizons, this model uses aggregated temporal features (e.g., seasonal averages), weather conditions (e.g., average temperature), spatial and demographic features (e.g., population density, income, POIs, transit ac-

cess, geofenced zones), and travel behavior indicators (e.g., average trip distances and usage distributions).

C.2 Analytical Capabilities and Decision-Support Tool

In support of scenario testing and tactical planning, a lightweight web-based interface was also developed. As in Chapter 5, the tool builds on a consistent hexagonal spatial framework and supports multiple layers of descriptive and predictive analysis. Planners can interactively explore demand patterns and contextual overlays using a dynamic map interface (Figure 1) or exploring temporal demand dynamics as well as configure daily and seasonal prediction models with relevant parameters (Figure 2). System-level plots also further summarize key metrics, including pickup/drop-off distributions, transit alignment, and spatial imbalance (Figure 3).



Figure 1: Main map interface showing pickup volumes by hexagonal cell overlaid with population background shading. The sidebar allows users to filter by year, season, demographic layer, POI, and operational zones.

The transit integration analysis in this study mirrors the methodology introduced in Chapter 5 for bike-sharing (Figure 4). In addition to spatial proximity, this study benefited from a valuable dataset provided by the SCOOTY team: boarding (“Ons”) and alighting (“Offs”) volumes at each transit stop. This enabled a more nuanced evaluation of potential first- and last-mile connections. Specifically, we compared: 1) *Transit Offs with e-scooter Pickups*, to identify areas where riders alight from transit and begin an e-scooter trip. 2) *Transit Ons with e-scooter*




Figure 2: Prediction configuration panels. (a) and (b) Daily and seasonal prediction inputs. (c) Time-lapse viewer for exploring temporal demand dynamics.




Figure 3: Various system-wide analytics interfaces. Visualizations include seasonal pickup/drop-off totals, demand distributions by zone, and transit stop-level performance indicators.

Drop-offs, to assess where users conclude their scooter trips and transition to transit.

This richer integration of demand volumes into spatial analysis allow to prioritize zones with high latent potential for multimodality. These capabilities provide spatially detailed, data-driven support for tactical micromobility planning and illustrate how the core contributions of this dissertation translate to other shared mobility systems.

C.3 Policy Implications and Future Directions

From a policy perspective, the ability to identify demand hotspots, evaluate multimodal integration potential, and simulate deployment scenarios offers practical value to cities and operators




Figure 4: Transit integration analysis based on: (a) transit stops, (b) transit hubs with high ridership volume.

aiming to optimize shared mobility performance. Future work may incorporate prescriptive optimization features, such as rebalancing strategies or parking incentive zones, and extend the models to reflect behavioral responses and network effects more explicitly. As data richness and spatial resolution improve, such frameworks will be increasingly vital to shaping sustainable, flexible, and equitable urban mobility ecosystems.

While this study is distinct from the dissertation's core chapters, it illustrates how shared methodological principles, particularly the integration of spatial analytics, predictive modeling, and interactive planning, can be extended to operational contexts and alternative micromobility modes. The findings in this appendix underscore the need for data-driven, locally responsive tools to support tactical planning, especially in hybrid and dockless systems where infrastructure is less fixed.