Titre: Title:	Data-Driven Methods for Inventory Management in Dock-Based Bike-Sharing Systems
Auteur: Author:	Maria Clara Martins Silva
Date:	2023
Type:	Mémoire ou thèse / Dissertation or Thesis
Référence: Citation:	Martins Silva, M. C. (2023). Data-Driven Methods for Inventory Management in Dock-Based Bike-Sharing Systems [Thèse de doctorat, Polytechnique Montréal]. PolyPublie. https://publications.polymtl.ca/57038/

POLYTECHNIQUE MONTRÉAL
affiliée à l'Université de Montréal

Data-driven Methods for Inventory Management in Dock-based Bike-sharing Systems

MARIA CLARA MARTINS SILVA
Département de génie informatique et génie logiciel

Thèse présentée en vue de l'obtention du diplôme de *Philosophiæ Doctor*
Génie informatique

Décembre 2023

POLYTECHNIQUE MONTRÉAL
affiliée à l'Université de Montréal

Cette thèse intitulée :

Data-driven Methods for Inventory Management in Dock-based Bike-sharing Systems

présentée par **Maria Clara MARTINS SILVA**
en vue de l'obtention du diplôme de *Philosophiæ Doctor*
a été dûment acceptée par le jury d'examen constitué de :

Guillaume-Alexandre BILODEAU, président
Daniel ALOISE, membre et directeur de recherche
Sanjay Dominik JENA, membre et codirecteur de recherche
Jonathan JALBERT, membre
Leandro C. COELHO, membre externe

DEDICATION

To Mack and Amora, who provided me with an infinite amount of love and support.

ACKNOWLEDGEMENTS

My most sincere thanks go to Professors Daniel Aloise and Sanjay Jena. I am deeply grateful for the opportunity to work under your supervision. I am truly grateful for the guidance and insightful feedback you provided me along this journey.

Mack and Amora, without you, this thesis would not exist. Your presence and support provided me with much-needed motivation and comfort during long hours of research and writing.

I extend my appreciation to my parents, João and Missilene, who gave me so much support since the beginning. I cannot imagine how difficult it must have been to see me grow from a distance but know that you will always be in my heart.

Thanks to my Canadian family, Ingrid, David, and Kendall, for bringing so much light into my life. Your warmth and hospitality have made my academic journey truly memorable.

To my friends, who I will not dare to list as I am afraid to miss one of the many people in my life, my deepest thanks. The countless joyful moments we shared were fundamental throughout this journey.

Finally, I would like to express my gratitude to professors Guillaume-Alexandre Bilodeau, Jonathan Jalbert and Leandro Coelho for accepting the invitation to compose the jury of my thesis.

RÉSUMÉ

Les systèmes de partage de vélos (BSS) offrent un moyen de transport écologique, atténuant la circulation et réduisant les émissions de gaz à effet de serre. De plus, ils favorisent un mode de vie plus sain pour leurs navetteurs. En raison de leurs nombreux avantages, les systèmes de partage de vélos ont connu une popularité croissante au fil des années. Néanmoins, satisfaire la demande de location et de retour peut s'avérer une tâche difficile car les stations ont une tendance naturelle à se déséquilibrer pendant la journée. Pour atténuer ce problème, les opérateurs BSS redistribuent les vélos entre les stations. Cependant, leur stratégie de rééquilibrage est souvent réactive, redistribuant les vélos une fois que les stations ont atteint le statut vide ou plein. Ce retard dans la résolution des déséquilibres des stations peut entraîner une perte de demande et une diminution de la satisfaction des clients. Pour cette raison, cette thèse propose trois nouveaux modèles basés sur les données qui contribuent à un rééquilibrage proactif dans le BSS.

Notre première contribution approfondit une stratégie tactique de rééquilibrage des opérations en suggérant des recommandations de rééquilibrage basées sur les données et triées en fonction de leur niveau d'urgence. À cette fin, nous proposons un modèle qui calcule automatiquement les intervalles d'inventaire et les valeurs d'inventaire cibles pour toutes les stations des systèmes. Basé sur ces mécanismes, notre modèle est capable d'identifier les gares avec un inventaire déséquilibré et d'établir l'inventaire optimal pour maximiser la proportion de trajets satisfaits dans les heures à venir. De plus, nous proposons quatre stratégies de priorisation qui trient les stations déséquilibrées en fonction de leur besoin d'intervention en fonction de différents critères. Après évaluation à l'aide de données de voyage réelles sur deux années contrastées, 2019 et 2020, l'une des stratégies a montré une réduction spectaculaire de la demande perdue jusqu'à 65%, et une autre a conduit à une diminution de 33% des opérations de rééquilibrage par rapport à la stratégie de rééquilibrage utilisée par un BSS du monde réel.

À mesure que les vélos électriques gagnent du terrain, de nombreux BSS les intègrent aux vélos ordinaires. Étant donné que les vélos électriques et classiques peuvent présenter des modèles de demande distincts, notre deuxième contribution explore les défis de rééquilibrage d'une telle intégration. Le nouveau modèle propose des recommandations de rééquilibrage basées sur des données, basées sur la demande prévue de vélos ordinaires et électriques. Le calcul proposé des intervalles d'inventaire et des valeurs d'inventaire cibles prend en compte la capacité totale de la station et le fait que tous les quais peuvent être utilisés par l'un ou

l'autre type de vélo. Pour évaluer les performances de notre modèle, nous simulons quatre cas différents dans lesquels les navetteurs adoptent des approches différentes lorsque le choix initial de vélo n'est pas disponible. Par rapport à la stratégie actuelle utilisée par un BSS réel, le modèle proposé a enregistré une baisse de 10% de la demande totale perdue et une baisse de 30% de la demande perdue pour les vélos électriques tout en maintenant un nombre similaire d'opérations de rééquilibrage moyen par heure.

Les modèles conventionnels qui résolvent le problème de repositionnement dynamique des vélos s'appuient souvent sur des intervalles d'inventaire calculés manuellement ou sur des valeurs d'inventaire cibles, ce qui peut conduire à des performances sous-optimales. Notre troisième contribution intègre donc les intervalles d'inventaire calculés automatiquement et les valeurs d'inventaire cibles dans un modèle d'optimisation de routage. Cela permet aux modèles proposés d'obtenir des itinéraires de rééquilibrage visant à atteindre un niveau satisfaisant de vélos et de quais disponibles pour toutes les stations du système. De plus, nous mettons en œuvre trois modèles de réoptimisation afin d'explorer les avantages de la mise à jour des inventaires de stations et les effets de la prise en compte d'horizons de planification plus ou moins longs. Des expériences informatiques sur des données synthétiques et réelles confirment que nos modèles surpassent systématiquement un modèle de routage traditionnel. Plus précisément, ils réduisent la perte de demande jusqu'à 38%, marquant une avancée significative dans l'optimisation des opérations BSS.

En tant que telle, cette thèse propose une exploration complète des principaux défis des opérations BSS, fournissant des informations exploitables et des méthodes innovantes qui peuvent considérablement améliorer le niveau de service des systèmes de partage de vélos.

ABSTRACT

Bike-sharing systems (BSS) provide an eco-friendly form of transportation, alleviating traffic and reducing greenhouse gas emissions. Furthermore, they promote a healthier lifestyle for their commuters. Due to their numerous advantages, bike-sharing systems have presented an increasing popularity over the years. Nevertheless, satisfying the rentals and returns demand can be a challenging task as the stations have a natural tendency to become unbalanced during the day. To mitigate this issue, BSS operators redistribute bikes among the stations. However, their rebalancing strategy is often reactive, redistributing bikes after the stations reach the empty or full status. This delay in addressing station imbalances can lead to lost demand and decreased customer satisfaction. For this reason, this thesis proposes three novel data-driven models that contribute to proactive rebalancing in BSS.

Our first contribution delves into a tactical strategy for rebalancing operations by suggesting data-driven rebalancing recommendations that are sorted according to their level of urgency. To this end, we propose a model that automatically computes inventory intervals and target inventory values for all stations in the systems. Based on these mechanisms, our model is able to identify the stations with unbalanced inventory and establish the optimal inventory to maximize the proportion of satisfied trips in the upcoming hours. In addition, we propose four prioritization strategies that sort the unbalanced stations according to their need for intervention based on different criteria. Upon evaluation using real-world trip data from two contrasting years, 2019 and 2020, one of the strategies showcased a dramatic reduction in lost demand by up to 65%, and another led to a 33% decrease in rebalancing operations in comparison with the rebalancing strategy used by a real-world BSS.

As electric bikes gain traction, many BSSs are integrating them alongside with regular bikes. Given that electric and regular bikes can present distinct demand patterns, our second contribution explores the rebalancing challenges of such integration. The novel model offers data-driven rebalancing recommendations based on the predicted demand for both regular and electric bikes. The proposed computation of inventory intervals and target inventory values considers the total capacity of the station and the fact that all docks can be used by either type of bike. To evaluate the performance of our model, we simulate four different cases in which the commuters adopt different approaches when the initial choice of bike is not available. Compared to the current strategy used by a real-world BSS, the proposed model registered a 10% drop in total lost demand and a 30% drop in electric bikes lost demand while maintaining a similar number of averaged rebalancing operations per hour.

The conventional models that address the Dynamic Bicycle Repositioning Problem often rely on manually computed inventory intervals or target inventory values, which may lead to sub-optimal performance. Our third contribution therefore integrates the data-driven automatically computed inventory intervals and target inventory values within a routing optimization model. This enables our proposed models to obtain rebalancing routes that aim to achieve a satisfactory level of available bikes and docks for all stations in the system. In addition, we implement three reoptimization models so we can explore the benefits of updating station inventories and the effects of considering shorter or longer planning horizons. Computational experiments on both synthetic and real-world data confirm that our models consistently outperform a traditional routing model. Specifically, they reduce lost demand by up to 38%, marking a significant leap in BSS operation optimization.

As such, this thesis offers a comprehensive exploration of key challenges in BSS operations, providing actionable insights and innovative methods that can greatly enhance the service level of bike-sharing systems.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF SYMBOLS AND ACRONYMS

BSS	Bike Sharing Systems
EBSS	Electric Bike Sharing Systems
GBT	Gradient Boosting Tree
MIP	Mix-Integer Program
DROB-I	Dynamic Rebalancing Optimization for BSS based on Inventory Intervals
DROB-T	Dynamic Rebalancing Optimization for BSS based on Target Inventory Value
DROB-LD	Dynamic Rebalancing Optimization for BSS minimizing Lost Demand

LIST OF APPENDICES

CHAPTER 1 INTRODUCTION

1.1 Thesis Motivation

Bike-sharing systems (BSS) are a shared mobility mode in which commuters have access to bikes scattered around the city. These bikes are available for rent for a specified period of time, enabling commuters to easily navigate urban areas without the need to own a personal bike. Its numerous advantages, such as easy access to bikes, sustainable transportation, and promotion of physical activity, has made BSSs a consolidated transportation alternative in different cities around the world. In 2022, there were roughly 1900 operational BSSs in over 1500 cities [87]. BIXI, a BSS located in Montreal, holds the title of the oldest large-scale BSS in North America, reaching more than 1.5 million trips in a single month in 2022. As more people embrace this mode of transportation, cities experience reduced carbon emissions and traffic conditions while overall increasing the quality of life[12, 85].

There are two main types of BSS: free-floating and dock-based. In the former, bikes can be parked at any public place that is within an allowed parking zone, as long as it follows the city's regulations. In the latter, bikes have to be plugged into docks at stations that have fixed locations and a limited dock capacity. In 2022, dock-based BSSs represented about 60% of the BSSs in the world [87]. This fraction is likely higher for large-scale BSSs since it was observed that free-floating BSSs can lead to congested streets or sidewalks, which can be a serious urban problem in a BSS with a high volume of bikes [19]. Also, bikes in free-floating systems are more likely to be vandalized, as observed in the first generation of BSSs [23]. In this study, the proposed rebalancing methods aim to improve the service performance of dock-based BSS.

Dock-based BSSs often suffer from inventory imbalance problems, caused by work-related user trips. It is typical for a significant portion of their users to commute from residential areas to commercial areas before work, and commute the inverse direction after work. The demand analysis presents in [80] reinforces this observation, highlighting that the peak hours occur before and after working hours, i.e. 8 am and 5 pm. Figure 1.1 presents the proportion of rentals (i.e., the number of rentals divided by the sum of rentals and returns) for each station averaged during July and August 2019 at 8 am (Figure 1.1a) and 5 pm (Figure 1.1b). Stations closer to 1 (blue) have a high demand for rental bikes, whereas stations closer to 0 (red) have a higher return demand at the specific time. The demand for the aforementioned hours appears to be inversely related, that is, regions experiencing high rental demand at 8 am tend to have high return demand at 5 pm, and vice versa. Based on the trip pattern

presented in Figure 1.1 we can identify the regions that are more likely to be residential and commercial. Nevertheless, having a wave of commuters renting or returning bikes at the same regions can lead to either empty or full stations and eventually result in lost rental or return demand. In the long run, unbalanced stations can affect the commuter's credibility in the BSS and, hence, decrease its user base.

In order to consistently provide both bikes and docks to their users, BSS operators coordinate bike redistribution among the stations, a process known as *rebalancing*. The logistics of the rebalancing can either be directly planned by the operators (operator-based rebalancing) or be planned indirectly through customer incentives (user-based rebalancing). Nowadays, large-scale BSSs often use both types of rebalancing to prevent inventory imbalance, such as BIXI (Montreal), Citi Bike (New York), and Divvy (Chicago). In this study, we specifically focus on operator-based rebalancing, which has shown to be more effective for dock-based BSS [23]. Also, the rebalancing decisions can be tailored without relying on user availability.

Coordinating rebalancing operations in a BSS poses a significant challenge, requiring careful consideration of various aspects. First of all, their rebalancing decisions have to take into account the demand for rentals and returns at a station level for the next hours. However, this is a complex task, given that many external factors (e.g., climatic conditions, day of the week, hour of the day, holidays, public events) can influence the demand [46, 119]. In Figure **??**, we can observe that even though we can identify regions with similar demand, there are also stations with high rental demand located near stations with high return demand. In the same line of thought, the demand for bikes has changed considerably in 2020 as the population adopted remote work due to the lockdown measures applied by the Canadian Government to fight the spread of COVID-19 [47, 108]. The abrupt change in demand makes it difficult to forecast demand under these circumstances.

In addition, several BSSs have expanded their offerings by incorporating electric bikes (also known as e-bikes) into their inventory, aiming to accommodate new preferences of commuters, while retaining the existing user base of regular bikes. BSSs that have both electric and regular bikes in their network, such as BIXI (Montréal), Citi Bike (New York), Santander Cycles (London), and Vélib' Métropole (Paris), are commonly referred to as *bimodal BSSs*. Compared to regular bikes, e-bikes have a higher speed, are easier to pedal, and provide more comfort for long-distance rides. On the other hand, regular bikes are cheaper, do not require battery charging, and attract commuters who look for health benefits. Therefore, providing both types of bike increases user adherence to the system, attracting higher revenues in the study presented in [134]. However, having two types of bikes in the BSS network adds a further level of complexity to demand forecasting, as the demand for each bike can be

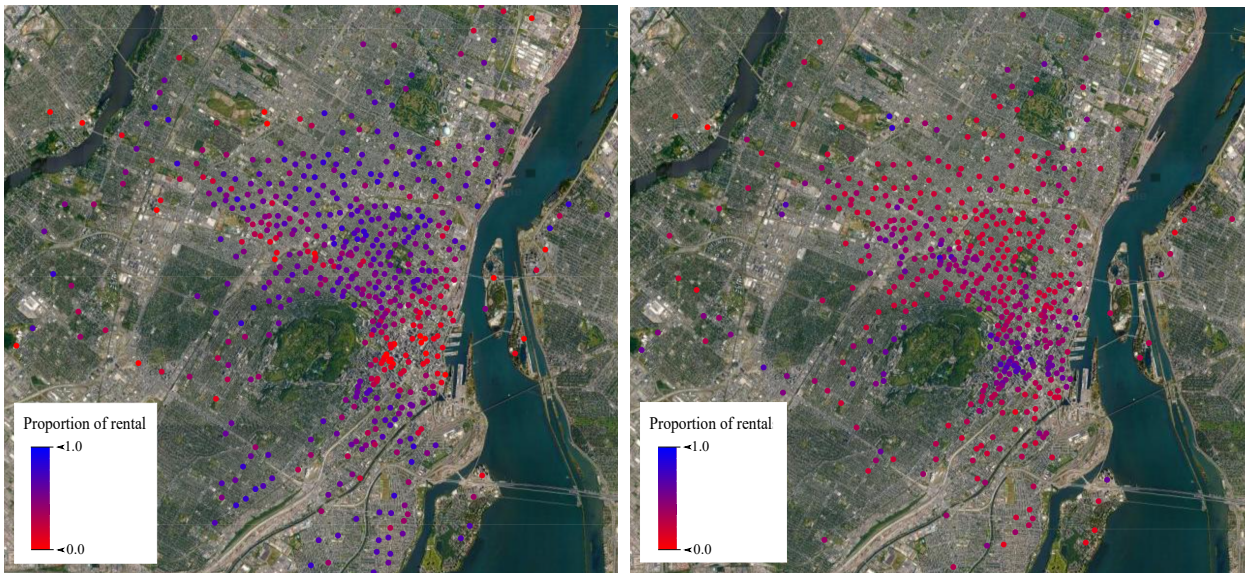

(a) 8 am (b) 5 pm

Figure 1.1 Rental demand over the total number of trips (rentals plus returns) averaged for July and August 2019 at BIXI.

considerably different depending on the region where the station is located. For example, stations located in flat areas are less likely to demand e-bikes than stations located in hilly areas. On the contrary, in regions with older and wealthier populations, it is expected that commuters will prioritize comfort over cost, resulting in a higher demand for e-bikes as compared to other areas.

Besides the demand prediction, operators have to take into account resource limitations when planning the rebalancing operations. In order to maintain the BSS's commitment to environmental sustainability and for economic reasons, it is crucial to avoid an extensive fleet of trucks redistributing all stations that deviate from their optimal inventory levels. The logistics of rebalancing also has to consider the costs of the fuel, rent of the trucks, the maintenance of the trucks, the driver's salary, etc. Therefore, operators need to strategically select a subset of stations for rebalancing, aiming to minimize future lost demand while aligning with the BSS's budget and sustainability objective.

To facilitate the rebalancing planning, operators prioritize bike redistribution among the stations classified as unbalanced. BSSs use different criteria to establish which stations are unbalanced. For instance, in NiceRide (Minneapolis, U.S) and Velo Antwerpen (Antwerp, Belgium) the operators consider a station unbalanced if their inventory is either completely empty or completely full [117, 122]. Meanwhile, in Vélo'v (Lyon, France) a station is unbal-

anced if the absolute difference between the number of arrivals and departures is larger than the standard deviation of the distribution of these values over all the stations [5]. In BIXI (Montréal, Canada) the operators use *inventory intervals* to define the unbalanced stations and *target values* to restore the station inventories after rebalancing. The inventory intervals represent an acceptable range in which the inventory can fluctuate while still satisfying a minimum expected demand. Thus, whenever a station's inventory lies outside its inventory interval, a rental or return alert is raised indicating the need to add or remove bikes at the unbalanced station. On the other hand, the target value represents the station's inventory that is able to maximize the satisfied trips. Nowadays, BIXI's inventory intervals and target inventory values are manually computed by their dispatching team.

Even though inventory intervals help operators to identify stations that do not meet the minimum service level for the next hours, the number of alerts commonly exceeds the rebalancing capacity. Figure 1.2 illustrates the average number of rental and return alerts per day for July, and August 2019 along with the maximum number of rebalancing operations during the peak hour for the same period. Thus, operators must select a subset of unbalanced stations to be rebalanced on a daily basis. These decisions are often made by their operators who use their expertise to estimate which stations are more likely to generate lost demand in the near future. Given the complexity, solely relying on human operators to predict the demand when analyzing the constraints of the rebalancing process, and based on it set the rebalancing schedule may result in a suboptimal outcome.

1.2 Research Objective and Contributions

In light of the challenges faced by BSS operators when it comes to devising effective rebalancing schedules, the major objective of this thesis is to propose novel data-driven strategies to improve the rebalancing decision-making process in BSS. More specifically, we aim to propose strategies that can provide rebalancing recommendations based on two mechanisms: inventory intervals and target inventory values. To ensure that these mechanisms can improve the service level in the following hours, they are computed considering the predicted rental and return demand. In addition, the computation of the inventory intervals and target inventory values must be easily adjusted to suit the rebalancing strategy followed by the operator. Thus, by introducing novel rebalancing strategies based on inventory intervals and target inventory values we aim to enhance the effectiveness of rebalancing operations and, hence, mitigate the lost demand caused by inventory imbalance. To achieve such goals, we propose three novel models able to provide rebalancing recommendations at a tactical and operational level.

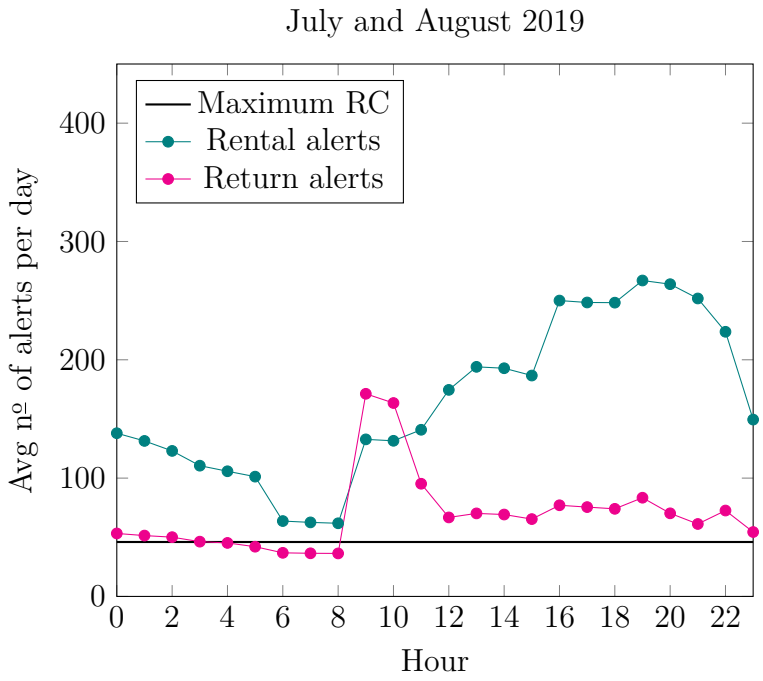

Figure 1.2 Maximum hourly rebalancing capacity and average daily number of unbalanced stations per hour at BIXI BSS during weekdays in July and August 2019 .

The major contributions of the proposed models are outlined below.

1. **Provide easy to implement rebalancing recommendations for BSSs.** First, we propose a mode to compute the inventory intervals and target values for dock-based BSS based on predicted demand – the future demand is estimated in accordance with the model presented in [53]. Notably, the integration of a robust predictive model assumes an even greater significance in scenarios of drastic demand fluctuations, as evidenced during the COVID-19 pandemic. In addition, the proposed model distinguishes from previous models by considering both rental and return service levels when setting the inventory interval for each station in the network.

   Then, we propose four different prioritization strategies that rank the unbalanced stations that are in most need of intervention. The prioritization strategies assist the BSS operators in selecting a subset of unbalanced stations when the number of such stations exceeds the BSS rebalancing capacity. Each strategy addresses a distinct aspect to enhance the BSS performance. The first strategy sorts the unbalanced stations based on their imminent risk of demand shortfall in the upcoming hours. The second strategy

emphasizes prioritizing stations based on the potential reduction in lost demand achievable through their immediate rebalancing. The third strategy prioritizes stations based on the deviation of their computed inventory intervals. Lastly, the fourth prioritization strategy considers geographical proximity, conferring higher priority to stations located closer to other stations in critical state. The results show that the prioritization strategies are able to reduce the estimated lost demand by 35% for the 2019 season data, and by 65% for the 2020 season data, as compared to BIXI's prioritization strategy.

2. **Automatically generate inventory intervals and target value for bimodal BSSs.** We propose a model to compute inventory intervals and target values for electric, and regular bikes in a dock-based EBSS. The inventory intervals and target inventory values are computed based on the predicted trips for each demand and the total number of docks. By considering the total number of docks, as well as their shared nature, our model manages to avoid unrealistic scenarios (e.g., the sum of the target values for each type of bike being greater than the total number of docks at the station) while still considering the demand for each type of bike. Thus, our model is able to provide individual rebalancing recommendations for BSS with both regular and electric bikes, bolstering the efficiency of their rebalancing operations.

   Furthermore, our model was designed to allow flexibility in configuring inventory intervals and target values, enabling different rebalancing strategies for each demand. Therefore, by knowing the user preferences, the inventory intervals and target values can be configured to meet their specific demands. The use of tailored inventory levels for each demand has been shown to be more effective than the strategy currently used by BIXI in most of the experiments, reducing the total lost demand by about 10% and the lost demand for e-bike by 30%, while maintaining an almost identical average hourly count of rebalancing operations. in one of the simulated scenarios. The results indicate that using a data-driven strategy to compute inventory intervals and target values can improve the performance of the BSS while respecting the maximum rebalancing capacity of the BSS.

3. **Incorporating inventory interval, and target values into routing optimizing model.** We investigate the benefits of using data-driven inventory intervals and target inventory values in the dynamic bicycle repositioning problem (DBRP). To this end, we propose two mixed-integer programming models, named DROB-I and DROB-T, that schedule rebalancing routes minimizing the deviation from inventory intervals and target inventory values, respectively. Both models are compared to a baseline optimization model that minimizes the expected lost demand. The demand prediction used in

the optimization models, as well as the creation of the inventory intervals and target inventory values, is estimated using a machine learning algorithm, specifically a Gradient Boosting Tree (GBT). For each model, we implement three reoptimization modes: static planning, rolling planning, and folding planning. These planning approaches provide valuable insights into the importance of updating the system status.

We evaluate the performance of the models through a simulator in which the inventory of the stations is estimated under the first-arrive-first-serve rule. The simulator accounts for the lost demand events across all stations throughout the entire planning horizon. The collected results show that the inclusion of inventory intervals and target inventory values brings significant benefits to BSS performance, particularly in scenarios in which the demand forecast is not highly accurate. This occurs because DROB-I and DROB-T aim to achieve a balance between available bikes and docks. In contrast, the baseline model focuses solely on minimizing demand loss and relies on trip prediction to identify the unmet demand for bikes or docks in the following hour. Thus, the baseline model runs the risk of ignoring stations that are either completely full or completely empty if the prediction fails to signal demand loss in the upcoming hour. The collected results show that DROB-T reduces lost demand by up to 34% and DROB-I manages to decrease lost demand by up to 28%, as compared to the baseline model. The experiments using real-world data confirm the outstanding performance of DROB-I over the basic dynamic model.

The secondary contributions of this work are enumerated below:

1. **Impact of the COVID-19 pandemic on BIXI's demand.** We conducted a comprehensive evaluation of inventory intervals and target inventory values using trip data from BIXI in 2019 and 2020. This study allows us to compare the demand patterns exhibited before and during the COVID-19 pandemic, thereby helping to understand the repercussions of the Canadian Government's lockdown measures. Our findings highlight not only quantitative differences in demand but also nuanced variations in demand behavior between these two distinct periods. For example, during the months of July and August 2020, the traditional peak in trip activity before and after regular working hours has notably dissipated, showcasing a clear response to the government's promotion of remote work during that period.

2. **Strategy to balance the number of bikes reallocated from the unbalanced stations.** The inventory intervals, combined with the prioritization strategies, are able to provide a sorted list of stations to be rebalanced according to their propensity to lose

future demand. However, they do not take into account the number of bikes reallocated from the stations during the rebalancing. To maintain a desirable balance where the number of bikes taken from congested stations aligns with the number of bikes added to less crowded stations, we devised a post-processing procedure. This procedure ensures a proportional adjustment in the number of bikes added and removed at the selected rebalancing stations, promoting effective rebalancing.

3. **BSS performance study in different user preference scenarios.** A study was developed using BIXI trip data in 2022 to simulate the performance of the BSS in different scenarios regarding commuters' preferences. The simulated scenarios were: (i) users never replace their desired bike with another type, (ii) users are flexible in their preferences and will always accept any available bike, (iii) users seeking regular bikes only accept an e-bike if the former is unavailable and (iv) users seeking for an e-bike only accept a regular bike if e-bikes are unavailable. We can then evaluate the performance of the inventory intervals in each users' preferences scenario.

4. **Performance analysis in different forecasting scenarios.** We investigate the performance of our proposed models, namely DROB-I and DROB-T, over different demand forecasting approaches. To this end, we applied two strategies to introduce noise into the trip predictions. The first one, the optimistic forecast, overestimates the trips while the second, the pessimistic forecast, underestimates the number of trips, in comparison with the original prediction. This allows us to grasp the significance of demand prediction and the impact of updating system status during the decision-making process of the rebalancing operations.

5. **Comparison of the performance of different reoptimization modes.** We applied different reoptimization modes in the optimization model: static planning, rolling planning, and folding planning. Static planning captures rebalancing strategies for each time period in the planning horizon at once. In rolling planning, the planning horizon is divided into stages, and at each stage, the rebalancing strategy is updated. Folding rolling planning is similar to rolling planning but it considers all the remaining time periods at each stage. Each reoptimization mode represents a different strategy in the logistics of the rebalancing and we evaluate the performance of these reoptimization modes in different scenarios.

6. **A trip data generator that is responsive to the weather.** We developed a data generator that is capable of emulating real-world trip data patterns. Thus, we can generate data for a fictional BSS with a downsized station network, thereby enhancing

the tractability of the optimization model. Different from other data generators presented in the literature, ours is able to create trip data that is responsive to different weather condition. We can therefore apply a more robust predictive model to estimate the future trips based on the climate conditions.

1.3 Thesis outline

The subsequent chapters of this document are structured in the following manner. Chapter 2 reviews the most relevant works in the area of demand prediction, tactical planning, and rebalancing planning. In the following chapters, three research articles are presented. Chapter 3 presents the proposed model that provides rebalancing recommendations based on demand prediction. Chapter 4 describes a novel model that computes inventory intervals and target values for BSS with two types of demand. In Chapter 5 we introduce two optimization models, named DROB-I and DROB-T, that incorporate inventory intervals and target inventory value into their objective functions. A comprehensive discussion of the proposed models is provided in Chapter 6. Finally, Chapter 7 summarizes the conclusions of this thesis.

CHAPTER 2 LITERATURE REVIEW

In this chapter, we review key studies and methods that address strategies to improve the performance of BSSs. More specifically, we delve into demand prediction and the rebalancing problem from a tactical and operational level. In the last section, we present a discussion to highlight the unique contributions of this thesis so as to position it within the existing literature.

2.1 Demand Prediction

Predicting the future rentals and returns in a BSS is an essential task for its effective. To this end, it is required to analyze historical trip data that inform about the number of rentals or returns that took place during specific time periods, such as hourly or daily intervals. This process helps uncover meaningful patterns and trends in bike usage. Figure 2.1 depicts an example in which raw trip data (left) is preprocessed and trips are aggregated into time periods of one hour (right). In addition, when processing the data, it is recommended to remove outliers, such as trips with a duration shorter than 1 minute or trips with missing values, to improve the learning process.

Raw data

Star_time	Start_station	End_time	End_station
13/09/2022 07:19:41	6023	13/09/2022 07:23:38	6614
13/09/2022 07:19:51	6212	13/09/2022 07:25:12	6203
⋮	⋮	⋮	⋮

Processed

Hour	Day	Month	Rental 6032	Return 6032	···
7	13	9	1	1	⋮
8	13	9	0	2	⋮
⋮	⋮	⋮	⋮	⋮	⋱

Figure 2.1 Example of raw trip data (left) and processed data (right)

In the literature, two primary approaches stand out for predicting demand in BSSs: the distribution-based models [2, 45, 68, 97, 106], and the regression models [6, 30, 50, 53, 56, 72, 76, 102, 127]. In distribution-based models, it is common to assume that the rentals and returns for station s follow a Poisson distribution with rates $\mu_s(t)$ and $\lambda_s(t)$, respectively. Both rates are estimated by averaging the historical rentals and returns (Maximum Likelihood Estimation for Poisson parameters). In [2], $\mu_s(t)$ and $\lambda_s(t)$ are estimated separately for each type of day (weekday, Saturday, and Sunday) to ensure that days with differing trip patterns

do not interfere with one another. Based on the estimated Poisson distribution, the authors compute the probability of rentals and returns for the next time period.

The distribution-based approach relies on the assumption that the trips follow a statistical distribution in all time periods. Regression models manage to predict the demand by identifying the relationship patterns between the trip demand and external features (time, weather, proximity to points of interest, etc), without assuming a specific trip distribution. Considering external features is an important advantage given that biking is an outdoor activity and, therefore, its demand is subject to the influences of the weather. Likewise, because it is a mode of transportation, bike-sharing is influenced by points of interest and temporal features. As such, studies have shows that algorithms that capture the impact of weather conditions and time features on the trip demand tend to provide more accurate prediction than predictions based on mean historical trip data [72, 74].

2.1.1 Regression Models

Regression models have the goal of predicting an output (also called dependent variable) based on a set of given features (called independent variables). In our context, the output is a continuous variable that represent the expected number of rentals and returns, while the features are external factors that influence the demand, such as time of day, day of the week, holiday, temperature, rain, humidity, proximity to a point of interest, etc. In order to learn the relation between features and the output, regression models iteratively refine their parameters and hyperparameters aiming to minimize a loss function. This function measures the difference between the predicted values and the ground truth, serving as a crucial guide that directs the algorithm towards improved prediction accuracy. Adopting the notation where y_i and $\hat{y}_i$ are, respectively, the ground truth and predicted value for the ith sample, Table 2.1 provides an overview of the most commonly used loss functions in machine learning.

Research papers on regression models for demand prediction in BSSs can be categorized according to the level of granularity chosen for the network, that is, the regression models can estimate the trips across the entire network [5, 46, 127], a cluster of stations [15, 56, 67, 68], or individual stations [6, 30, 53, 76, 102, 125]. Given that the primary focus of this thesis revolves around rebalancing recommendations at the station level, our attention is directed toward the latter category. In this case, regression models are designed to predict upcoming demand at each station, treating rentals and returns as distinct demands. Thus, to put this into perspective, in a network with 50 stations, one requires 100 regression models, two per station, to predict the rentals and return of the entire network. Figure 2.2 illustrates

Table 2.1 Loss functions

Loss function	Equation
Mean Square Error (MSE)	$\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2$
Root-Mean-Square Error (RMSE)	$\sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}$
Mean Absolute Error (MAE)	$\frac{1}{n}\sum_{i=1}^{n}\lvert y_i - \hat{y}_i\rvert$
Cross Entropy Loss	$-(y_i log(\hat{y}_i) + (1 - y_i) log(1 - \hat{y}_i))$

a regression model tailored to forecasting rentals at station 601, taking into account hour, day, temperature, and rain information. Each data sample is structured such that each one contains detailed hourly information regarding these features, leading to corresponding hourly rental predictions.

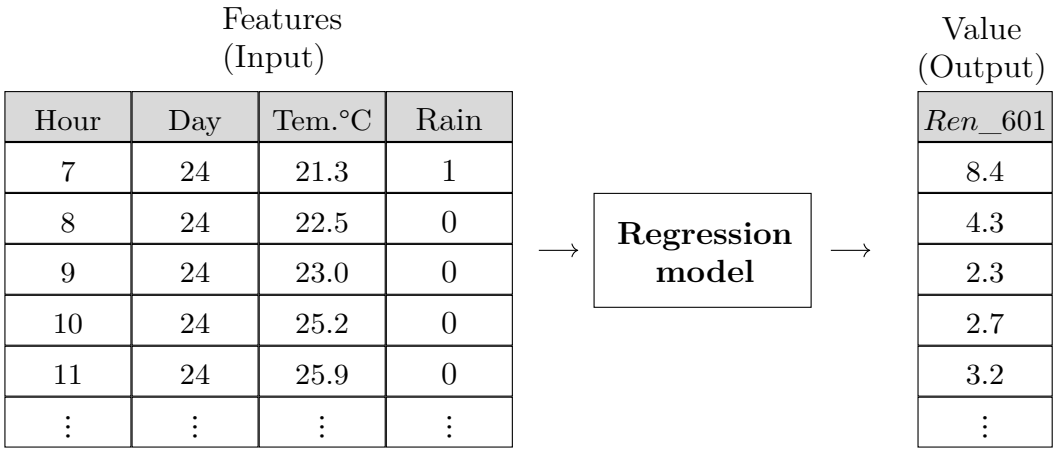

Figure 2.2 Example of a regression for predicting the number of bike rentals at station #601.

Most regression models employ weather and temporal data as input features [50, 53, 72, 127]. Nevertheless, numerous studies highlight the influence of other feature categories on bike demand. For example, [5] includes the number of available bikes and the total number of commuters with a membership of the BSS in the regression model. [67, 74, 76, 83] use network information, such as the capacity of the stations, location of the stations, and real-time trip data, as input to their regression model. The work presented in [102] employs features for each station to indicate if the three closest stations, to the station in question, are either empty or full – they seek to capture the demand occasioned by missing bikes or docks in the nearby stations. Geographic aspects, such as proximity to points of interest

(e.g., universities, restaurants, malls, grocery stores, and metro stations) and land use, are explored in [30, 39, 76]. Demographic features, like population density, employment density, and population zones, find relevance in regression models, as demonstrated in [25, 46]. These diverse feature categories enrich the understanding of trip demand patterns in BSS, enabling more accurate and comprehensive predictions.

There exists a wide variety of machine learning regression models. Among the works addressing the rebalancing problem in BSS, the most recurrent algorithms to appear in the literature are: linear regression [5, 30, 96, 100, 102, 126], random forests [53, 74, 96, 103, 127], gradient-boosting trees [6, 53, 56, 74, 96, 125, 127], and neural networks [16, 76, 83, 103, 125]. Works comparing the accuracy of trip prediction in BSS using machine learning algorithms have consistently shown that tree-based predictive models outperform other algorithms [6, 53, 127]. Given this superior performance, the models presented in Chapters 4, 5, and 6 of this work also employ tree-based predictive models. In the following section, we will delve deeper into the concepts of random forests and gradient-boosting trees.

2.1.2 Tree-based predictive models

Random forests are part of a class of algorithms known as *Ensemble Methods*, which are characterized by the combination of multiple models to improve the prediction process. The fundamental idea behind random forest is to create a set of decision trees [78], each of which is grown independently and exhibits random variations. These variations are introduced through a technique known as *bagging* that generates new datasets (D') by randomly choosing elements from the original dataset (D) with replacements. As a result, some elements from D can be duplicated in D', whereas other elements might be absent. Consequently, each tree in the forest receives a different set of data, which contributes to model diversity. In a classification problem, the final prediction is determined by the majority classification of the decision trees, while in a regression problem, the final result is the average between the results from each tree. The visual representation of the random forest is presented in Figure 2.3, which illustrates a scenario with n decision trees, each trained with unique datasets.

The gradient boosting tree (GBT) also fits into the category of Ensemble Methods as it benefits from the combination of multiple decision trees. However, different from the Random Forest that trains all the decision trees simultaneously, GBT adds decision trees sequentially so that each new tree corrects the errors of the previous one and improves the overall prediction – method of training known as *boosting*. To gauge the accuracy of each new prediction, GBT employs loss functions that quantify the disparity between predicted values and actual ground truth. This iterative, corrective process makes GBT a powerful tool for predictive

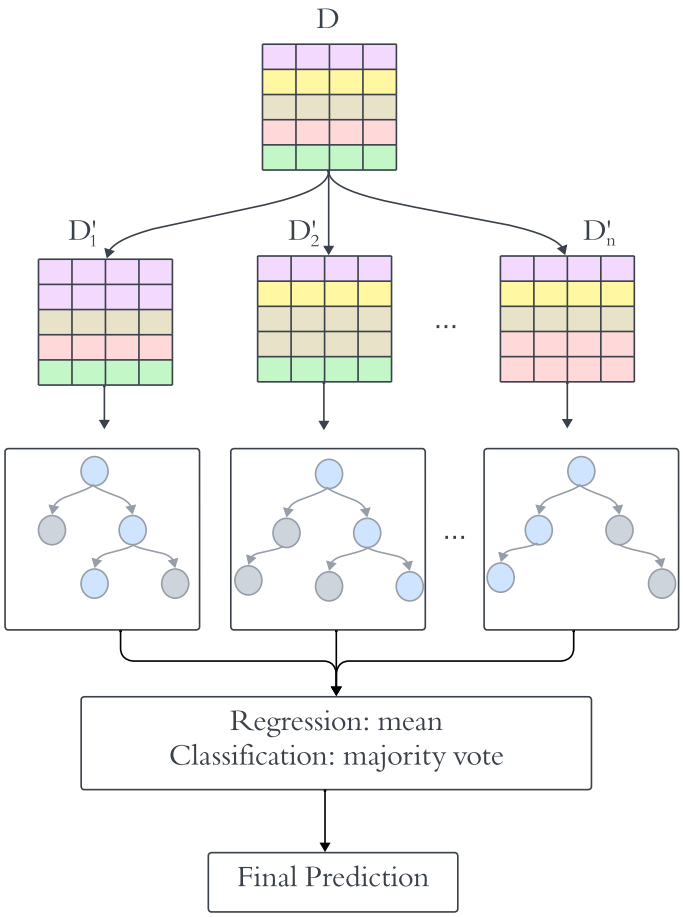

Figure 2.3 Example of a random forest with n decision trees

modeling.

To facilitate our explanation, we employ the following notation:

- i represents the index of the dataset, ranging from 1 to n;
- x_i denotes the input to the predictive function;
- y_i signifies the observed value;
- $\hat{y}_i$ represents the predicted value;
- F is the predictive function, and thus, $F(x_i) = \hat{y}_i$;
- L stands for the loss function, which measures the distance between y and $\hat{y}$.

The algorithm starts by assuming that the prediction of all samples is the average of y so the first predictive model is $F_1(x_i) = \overline{y}_i$. Then, the algorithm computes the gradient of $L(y_i, \hat{y}_i)$ with respect to $\hat{y}_i$. The gradient of the loss function indicates the direction in which each predicted value $\hat{y}_i$ should be adjusted to minimize the loss function. In the following step, a decision tree model is trained aiming to predict the gradient values rather than the output y. This decision tree, named F_2, captures the relationship between the input data x and the gradient of the loss function. This new decision tree is added to the model with a learning rate γ_2 that is computed as:

$$\gamma_2 = \underset{\gamma}{\operatorname{argmin}} \left[\sum_{i=1}^{n} L(y_i, F(x_i) + \gamma F_2(x_i)) \right]. \tag{2.1}$$

Resulting in a model:

$$F(x) = F_1(x) + \gamma_2 F_2(x). \tag{2.2}$$

The model continues to add decision trees, following the steps previously described until a specific threshold is reached.

2.2 Tactical Planning

According to [118], rebalancing on the tactical level consists of determining inventory levels at individual stations to compensate for the demand fluctuations along the day. These inventory levels can be either the ideal number of bikes a station should have (target inventory value) or an acceptable range for the number of docked bikes (inventory intervals) in a station [21]. Thus, when the rental demand is significantly higher than the return demand, the inventory levels tend to be situated closer to the maximum capacity of the stations. This adjustment ensures that a sufficient number of bikes is readily available to meet the high rental requirements. Conversely, in a scenario where return demand surpasses rental demand, the inventory levels should shift closer to zero. This aligns with the need to provide docks and prevent stations from becoming overstocked. Given the nuances involved in this process, manual computation of inventory intervals and target inventory values, although a common practice, often results in sub-optimal rebalancing performance. Some works have addressed this problem and proposed solutions to automate this decision-making process considering the future demand and the network configuration (e.g., the number of stations and their capacities).

2.2.1 Inventory Model

A common approach in the literature is to model the trip demand as a Markov chain. In this context, each potential inventory level for station s is denoted as a distinct state, with transition probabilities between states determined by the rates of rentals (μ_s) and returns (λ_s). This approach operates under the assumptions that: (i) rentals and returns are random processes with fixed rates in a given time period; and (ii) the trips are independent events and, therefore, future trips are not affected by past trips. For a visual representation of this concept, Figure 2.4 depicts the Markov chain that represents the dynamics of the bike inventory level for station s, with maximum capacity C, at time period t.

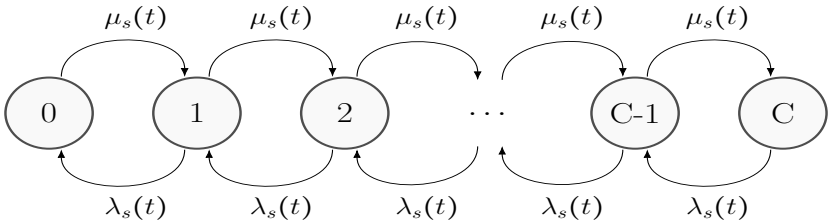

Figure 2.4 Markov chain that represents the dynamics of the stations inventory.

This Markov chain can be conceptually interpreted as an $M_1/M_2/1/K$ queue in which each parameter holds a specific significance in characterizing the system's behavior. M_1 and M_2 represent the time between the arrivals and returns at a station. Assuming that the rentals and returns follow the Poisson distribution, the intervals between rental and returns follow an Exponential distribution with same rate. The third and fourth parameters represent the number of servers, with each station serving as one server, and the maximum capacity of each server, that is the total number of docks available at the station.

2.2.2 Target Inventory Value

The target inventory value is computed aiming to find the initial inventory level for a given time period that optimizes a performance function, which serves as a metric for evaluating the BSS's performance. A variety of strategies have emerged to compute target inventory values, each tailored to address different facets of BSS optimization. For example, in [97], the target inventory represents the initial inventory that is able to minimize the User Dissatisfaction Function – which measures the performance of the stations according to the potential lost rentals and returns. In [21], the target inventory is computed as the initial inventory that minimizes the Journey Dissatisfaction Function. This function measures the satisfaction of the commuter during the whole process of renting and returning a bike, taking into

consideration waiting times and station switches. The authors in [72] estimate the target inventory value as the initial inventory that is able to maximize the time window during which the station is considered balanced. In the study presented in [45], the authors compute the probabilities associated with a station either becoming empty or reaching full capacity given an initial inventory. Then, the target inventory is selected as the inventory value for which the probabilities of transitioning into either an empty or full status are equal. The models previously mentioned proposed target inventory values that prioritize rentals and returns equally. In [50], the authors establish the target inventory value for central stations, i.e., high-demand stations within a cluster of stations, by estimating the initial inventory that minimizes the likelihood of a station reaching full or empty status. They further refine their approach by weighting the importance of rentals and returns at each central station, using the proportion of rentals and returns in the cluster to which the central station belongs. In [53], the target inventory value represents the initial inventory that is able to maximize the proportion of satisfied trips for the considered time period. The authors introduced a hyperparameter, named α, that operators can manually adjust according to their preference to satisfy rental or returns.

2.2.3 Inventory Interval

Regarding the computation of a inventory interval, a critical factor to consider is the gap between its upper and lower bounds. Modifying the gap size can directly influence the rebalancing strategy. Narrow inventory intervals restrict a station's inventory within a small range, which tends to trigger more rebalancing operations due to the limited buffer. On the other hand, wider inventory intervals provide a large inventory range, potentially reducing the requirement of frequent rebalancing operations but leaving the station more susceptive to lost demand cases. In [7], the authors propose a simplistic strategy in which inventory intervals are a fixed percentage (β) of the station's total capacity. So, for a station s with 20 docks and β = 0.1, the lower bound is 2 and the upper bound is 18. Different values of β were tested and the results indicate that higher values of β (narrower inventory intervals) consistently triggered more rebalancing operations and increased the number of served stations. However, it was observed that a higher number of rebalancing operations did not necessarily correlate with a reduction in lost demand. The value of β that achieved the best rebalancing performance was 0.2. The work present in [106] computes the inventory intervals as the range of the initial inventory that is able to achieve a minimum service level for a time period, that is, the inventory intervals are the range of the initial inventory that can guarantee that a certain proportion of the total demand will be satisfied. The authors carried out experiments with inventory intervals that guarantee a service level of 85%. In this case, their approach to

estimating the trip prediction, Maximum Likelihood Estimation, disregards the influence of the weather on the trip demand. The authors of [53] propose an extension of the previous work by adding a hyperparameter to control the gap between the upper and lower bound. This hyperparameter can assist the operator to adjust the inventory intervals according to their desired rebalancing strategy. Also, differently from the results obtained by [7], the authors observed that narrower inventory intervals always resulted in fewer cases of lost demand than wider inventory intervals.

Addressing the rebalancing problem from a tactical perspective can lead to significant enhancements in the rebalancing process. The studies previously discussed provide valuable insights for operators, assisting them in identifying which stations require rebalancing and the ideal number of bikes to be positioned in each one of them. Even more importantly, by predicting in advance those ideal numbers for upcoming hours the operators can adopt a preventive rebalancing approach, redistributing bikes across the network in order to meet the future demand. However, exclusive tactical-level decisions are often based on strong assumptions. Many of these models may overlook critical factors such as the total number of bikes in the network and the specific rebalancing constraints inherent to BSSs in practice. As such, in order to address such limitations, an operational decision-making view of the rebalancing problem is required.

2.3 Operational Planning

At the operational level, the rebalancing problem addresses the redistribution of bikes among stations to provide a satisfactory service level in the system [118]. The rebalancing problem can be divided between static rebalancing [2, 10, 13, 22, 27, 84, 98, 106, 118] and dynamic rebalancing [20, 43, 63, 70, 73, 117, 123]. Static rebalancing (also called overnight rebalancing) assumes that the stations' inventories remain unaltered during the rebalancing process, being often related to BSSs that opt to rebalance their stations at night and, hence, avoid traffic jams and parking problems [98]. Dynamic rebalancing assumes that the stations' inventories can change during the rebalancing process, taking into account the inventory fluctuations and expected user trip demand in the decision-making process. The dynamic approach is usually a better approximation of reality, since most of the BSS provide service uninterruptedly, and the reallocation of the bikes often occurs simultaneously with rentals and returns. For this reason, we will delve deeper into dynamic rebalancing problems.

Assuming that rentals and returns occur simultaneously to rebalancing operations introduces new challenges to the rebalancing problem. Given that rental and returns are aggregated within time periods, rather than being timestamped individually, the optimization models

can make assumptions about the chronological order of the trips and the rebalancing. For example, for each time period, the model can either consider that rentals, returns and rebalancing operations occur consecutively or that all the events occur simultaneously. The order of the events can directly affect the rebalancing strategies as it influences the availability of bikes and docks. A model that assumes that rentals occur before returns and rebalancing is more likely to estimate a higher number of rental lost demand and, hence, add more bikes to the stations – with the opposite logic if returns precede other events. On the other hand, a model that assumes that all events occur simultaneously can neglect stations with high demand if the demand for rentals and returns is similar. The majority of the works in the literature assume that these events happen simultaneously [20, 43, 63, 70, 117, 123], but a few works proposed a sequential approach [42, 73].

Another aspect that must be considered is the limitated number of trucks in the system as well as the number of stations each truck can rebalance in one time period. This adds another layer of complexity to the problem, as it necessitates careful allocation of resources to ensure that the most critical stations receive attention and that the limited fleet is optimally used. Works assuming a single truck are usually limited to small-scale BSSs or to a cluster of stations [13, 22, 106]. By considering multiple trucks, rebalancing models can relocate more bikes between the stations, which is beneficial for BSS with a considerable number of stations [20, 63, 73]. Some rebalancing models allow trucks to rebalance multiple stations within a single time period [10, 14, 27], while others restrict them to visit one station per period [22, 26, 70, 98].

The dynamic rebalancing problem has been widely investigated in the literature and different approaches have been proposed aiming to optimize the rebalancing decision-making process. Most optimization models in this area formulated the problem as a mixed-integer linear programming (MIP) models [see, e.g. 10, 13, 20, 22, 27, 63, 70, 73, 84, 98, 106]. MIP offers a structured framework that allows for the mathematical representation of the system's features. Thus, the rebalancing routes can be drawn aiming to optimize an objective function (e.g., minimizing lost demand, minimizing rebalancing cost, minimizing rebalancing time, etc.). By adding constraints to the model, it is possible to establish the operational boundaries of the system (e.g., bound the inventory per station, set a maximum route per truck, ensure a maximum number of bikes in each truck, etc.). Additionally, the models can receive a set of inputs with information about the network (e.g., number of stations, number of trucks available and their respective capacities, predicted future demand, distance between stations, etc.). This way, an optimization solver applied to the mathematical model can compute the best outcome for the decision variables of the formulated problem, i.e., an optimization solver can find the optimal route for each truck at each time period to ensure

that the objective function achieves its best possible outcome.

To better illustrate an optimization model formulated as an MIP, we present the work introduced in [70] that addresses the dynamic rebalancing problem. The decision variables and input notations are detailed in Table 2.2. The inputs of this model are constant values that characterize the BSS. Note that the authors included the inputs $f_s^{+,t}$ and $f_s^{-,t}$ so that the model can set rebalancing schedules considering the predicted demand. The decision variables d_s^t, $x_s^{+,t}$ and $x_s^{-,t}$ model the inventory of the station. On a similar note, $\hat{d}_v^t$ varies according to the pickup and drop out carried out by each vehicle. The decision variables $r_{s,v}^{+,t}$, $r_{s,v}^{-,t}$ and $z_{s,v}^t$ express the rebalancing and the routing decisions made by the model.

Table 2.2 Definition of parameters

Type	Notation	Definition
Input	S	The set of stations
	V	The set of vehicles
	T	The set of discreted time-periods
	C_s	The capacity of station $s, s \in S$
	$\hat{C}_v$	The capacity of vehicle $v, v \in V$
	L_t	The duration (in minutes) of time period $t, t \in T$
	d_s^1	The initial number of bikes in the station $s, s \in S$
	$\hat{d}_v^1$	The initial number of bikes in the vehicle $v, v \in V$
	$z_{s,v}^1$	The initial location of each vehicle $v, s \in S, v \in V$
	$f_s^{+,t}$	The expected rental demand at station s in period t
	$f_s^{-,t}$	The expected return demand at station s in period t
Decision Variables	d_s^t	The number of bikes available in station $s \in S$ at the beginning of period $t \in T$
	$\hat{d}_v^t$	The number of bikes in vehicle $v \in V$ at the beginning of period $t \in T$
	$x_s^{+,t}$	The number of successful bike trips starting from station $s \in S$ in period $t \in T$
	$x_s^{-,t}$	The number of successful bike trips ending at station $s \in S$ in period $t \in T$
	$r_{s,v}^{+,t}$	The number of bikes picked up at station $s \in S$ by vehicle $v \in V$ in period $t \in T$
	$r_{s,v}^{-,t}$	The number of bikes dropped off at station $s \in S$ by vehicle $v \in V$ in period $t \in T$
	$z_{s,v}^t$	$z_{s,v}^t = 1$, if vehicle $v \in V$ is at station s in period $t \in T$; 0 otherwise

$$min \quad \sum_{s\in S}\sum_{t\in T}(f_s^{+,t} - x_s^{+,t}) + \sum_{s\in S}\sum_{t\in T}(f_s^{-,t} - x_s^{-,t}) \tag{2.3}$$

$$\text{s.t.} \quad \hat{d}_v^{t+1} = \hat{d}_v^t + \sum_{s\in S}(r_{s,v}^{+,t} - r_{s,v}^{-,t}) \qquad \forall\ v \in V,\ t \in T \tag{2.4}$$

$$d_s^{t+1} = d_s^t - \sum_{v\in V}(r_{s,v}^{+,t} - r_{s,v}^{-,t}) - x_s^{+,t} + x_s^{-,t} \qquad \forall\ s \in S,\ t \in T \tag{2.5}$$

$$\sum_{s\in S} z_{s,v}^t = 1 \qquad \forall\ v \in V,\ t \in T \tag{2.6}$$

$$r_{s,v}^{+,t} + r_{s,v}^{-,t} \le \hat{C}_v z_{s,v}^t \qquad \forall\ s \in S,\ v \in V,\ t \in T \tag{2.7}$$

$$0 \le \hat{d}_v^t \le \hat{C}_v, 0 \le d_s^t \le C_s \qquad \forall\ s \in S,\ t \in T \tag{2.8}$$

$$0 \le x_s^{+,t} \le f_s^{+,t}, 0 \le x_s^{-,t} \le f_s^{-,t} \qquad \forall\ s \in S,\ t \in T \tag{2.9}$$

$$0 \le r_{s,v}^{+,t}, r_{s,v}^{-,t} \le \hat{C}_v \qquad \forall\ s \in S,\ v \in V,\ t \in T \tag{2.10}$$

$$z_{s,v}^t \in \{0,1\} \qquad \forall\ s \in S,\ v \in V,\ t \in T \tag{2.11}$$

Objective function (2.3) minimizes the difference between the expected trips ($f_s^{+,t}$ and $f_s^{-,t}$) and the observed trips ($x_s^{+,t}$ and $x_s^{-,t}$). In other words, this model sets a rebalancing plan that minimizes the number of lost rental and return demands. Constraints (2.4) update the number of bikes in vehicle v and time period $t+1$ by summing the number of bikes and the picked up/dropped off operations in t. Constraints (2.5) update the number of bikes at station s and time period t by summing the current inventory, the number of bikes added or removed by the truck and the successful trips. Constraints (2.6) limit each truck to visit a single station per time period. Constraints (2.7) establish that the truck can only rebalance a station if it is located at the stations. Constraints (2.8) bound the inventory of s and v to 0 and their respective maximum capacity. The following constraints (2.9) ensure that the number of successful trips is bounded by the predicted rentals and returns. Constraints (2.10) limit the number of pickups and drop offs to the maximum capacity of each truck. The final constraints (2.11) specify that $z_{s,v}^t$ are binary decision variables.

The optimization model presented in Liang et al. [70] obtains rebalancing routes based on minimizing the expected lost demand, a common rebalancing strategy found among studies addressing this topic [2, 20, 42, 73]. Other objective functions present in the literature are: minimizing rebalancing costs [13, 22, 27, 118], minimizing rebalancing time [4, 106], minimizing deviation from target inventory values [45, 62], minimizing deviation from inventory intervals [120], and minimizing two or more metrics at the same time [86, 98].

Regarding the constraints, [42] presents a MIP formulation in which the trucks can visit multiple stations in a single period but the total duration of the rebalancing is bounded by the duration of the period. In [13], the authors include a set of constraints to ensure that the stations reach the target inventory value at the end of the rebalancing. Similarly, [106] add constraints so that all the stations reach a balanced status, i.e., their inventory is within

the inventory interval by the end of the rebalancing. The works [32], [26] and [10] include constraints that ensure that the trucks begin and end the route at a bike depot.

Unfortunately, most models addressing the rebalancing problem formulated as MIPs face limitations in applicability, particularly in the context of large-scale BSS, due to their high computational complexity. Some works have proposed new strategies to mitigate this problem by reducing the size of the problems solved. [63, 106] employ a *Cluster-First Route-Second* strategy in which the stations are clustered based on their characteristics, being each cluster assigned to one or multiple trucks. In this case, they assume that each cluster is 'self-sufficient', that is, stations can be balanced by redistributing bikes within a cluster. To further improve tractability and runtime, [20] utilize two decomposition methods: Dantzig-Wolfe decomposition and Benders decomposition. These methods aid in breaking down the problem into more manageable components. [8] propose a temporal and spatial decomposition that divides the rebalancing scheduling into subproblems, where each subproblem is solved individually for each vehicle considering a single time period and a subset of the stations. The models presented in [117] and [112] follow a request-based approach in which the rebalancing routes are computed based on a list of requests that inform which stations should be rebalanced and how many bikes need to be added or removed in these stations. In their models, the rebalancing routes are computed considering only the subset of stations with assigned requests, therefore, narrowing the scope of the problem and improving practicality. Nevertheless, the authors assume that these requests are given as input to their model.

2.4 Discussion

Here we focus on gaps in the existing research about inventory management and rebalancing operations. We highlight three main points that result in the proposed themes in this thesis. By critically examining these points, we aim to underscore the need for data-driven methodologies to enhance the efficiency and effectiveness of BSS operations.

2.4.1 First Theme

There are numerous studies addressing inventory management in BSS. However, it is noteworthy that a significant portion of these works do not employ robust mechanisms to estimate future demand, disregarding the effects of weather and festivities on the demand. Furthermore, most published papers address inventory intervals and target inventory values separately. In Hulot et al. [53], the authors propose a model that resembles the strategy used by BIXI, in which both inventory intervals and target inventory values are combined to

provide rebalancing recommendations – indicating which stations are unbalanced and how many bikes have to be added or removed from these stations. However, they assume that all stations that raised an alert can be rebalanced within an hour. As shown in Figure 1.2, the number of alerts, hence, unbalanced stations, is often greater than BIXI's maximum rebalancing capacity. Therefore, in a real scenario, operators often have to select a subset of stations to be rebalanced based on their expertise, which likely results in a sub-optimal performance. In [117] and [112], it is proposed that the rebalancing recommendations, named as request, receive a priority weight according to their need for rebalancing. However, it is assumed that the requests, as well as their priority, are given. Therefore, the use of data-driven rebalancing recommendations with a priority weight automatically computed remains unexplored.

2.4.2 Second Theme

E-bikes are a recent trend in BSS and, hence, the number of works addressing this new commodity is still quite limited. To the best of our knowledge, existing works that focus on inventory intervals and target inventory values primarily consider a single commodity in the system. Such models do not suit bimodal BSSs, as they can generate unrealistic outcomes, e.g., the target of both bike types surpassing the total number of docks, or a full station failing to be identified as unbalanced since neither demand has exceeded its upper or lower bounds. Thus, to properly compute the inventory intervals and target inventory value for a dock-based bimodal BSS, the model needs to consider the maximum capacity of the station and the fact that docks can be used by both types of bikes. In the literature, the works addressing bimodal BSSs remain in the scope of strategic planning, that is, the proposed models aim to improve the configurations of the network, such as the number and location of the stations, number of bikes in the network, and number of trucks [118]. For example, [134] proposes a model to find the optimal number of mechanical and electric bikes in the network, considering the ride time between stations, the rental fee, and fatigue levels. [79] presents a model to find the location of the stations in the city of Lisbon and the number of bikes in the network that maximizes the BSS revenue. Thus, there is a gap in the literature of works addressing the automatic computation of inventory intervals and target inventory values for bimodal BSS.

2.4.3 Third Theme

In the domain of Dynamic Bicycle Repositioning Problem (DBRP), the models focusing on improving service satisfaction usually aim to minimize the expected lost demand [2, 20, 42, 73]

or to minimize the deviation from the target inventory values [24, 29, 91]. The estimation of lost demand is a tricks task as it concerns unobserved data, that is, there is no record of how many commuters fail to rent or return a bike at a specific station due to the lack of bikes or docks. Training a predictive model using historical trips then results in a model that essentially ignores unobserved unsatisfied trips. As such, rebalancing stations based on such predictions is essentially aiming at achieving what has already been achieved in the past. The alternative rebalancing strategy, minimizing the deviations from the target inventory values or inventory intervals, is a more cautious approach that ensures that stations maintain an adequate number of bikes and docks at all times. As most published papers restrain their work to a single rebalancing strategy, comparing the performance of different objective functions, while using a robust predictive model to predict the demand, can bring an important contribution to the literature. Furthermore, most existing works align with a strategy akin to BIXI's, using manually computed target inventory values, which are likely to be sub-optimal. Thus, the advantages of incorporating data-driven inventory intervals and target inventory values into rebalancing optimization need to be investigated.

CHAPTER 3 ARTICLE 1: DATA-DRIVEN PRIORITIZATION STRATEGIES FOR INVENTORY REBALANCING IN BIKE-SHARING SYSTEMS

Authors: Maria Clara Martins Silva, Daniel Aloise e Sanjay Dominik Jena.

Maria Clara Martins Silva: Development and implementation of the proposed model and simulation, experiments, validation, trip data analysis, literature review, writing and results analysis. **Daniel Aloise:** Development of the proposed model, writing, results analysis and supervision. **Sanjay Dominik Jena:** Development of the proposed model, writing, literature review, results analysis and supervision.

Under first round of revision at *OMEGA* [1]

Date: September 13th, 2023

Abstract. The popularity of bike-sharing systems has constantly increased throughout the last years. Most of such success can be attributed to their multiple benefits, such as user convenience, low usage costs, health benefits and their contribution to environmental relief. However, satisfying all user demands remains a challenge, given that the inventories of bike-sharing stations tend to be unbalanced over time. Bike-sharing system operators must therefore intervene to rebalance station inventories to provide both available bikes and empty docks to the commuters. Due to limited rebalancing resources, the number of stations to be rebalanced often exceeds the system's rebalancing capacity, especially close to peak hours. As a consequence, operators are forced to manually select a subset of stations that should be prioritized for rebalancing. While most of the literature has concentrated either on predicting optimal station inventories or on the rebalancing itself, the identification of critical stations that should be prioritized for rebalancing has received little attention. Given the importance of this step in current operating practices, we here propose four strategies to select the stations that should be prioritized for rebalancing, using features such as the predicted trip demand, as well as the inventory levels at the stations themselves and their surrounding stations. We evaluate the performance of these prioritization strategies by simulating real-world trips using data from 2019 and 2020, each of which exhibits distinct travel patterns given the restrictive measures implemented in 2020 to prevent the spread of COVID-19. One of these strategies significantly improves the system's performance by reducing the lost demand by up to 65%, while another strategy reduces the number of required rebalancing operations by up to 33%, when compared to the prioritization scheme currently used within our bike-sharing system

[1] available at [80]

use case. Finally, one prioritization strategy encourages the selection of stations that are geographically clustered, which may facilitate rebalancing operations afterward.

Keywords. *Bike-sharing, Demand prediction, Rebalancing, Inventory management.*

3.1 Introduction

Demand for bike-sharing systems (BSS) has constantly increased over the recent years, as they continue to provide various advantages: they are typically simple to use and do not require previous reservation; they have been shown to be an environmentally friendly transportation mode by reducing the ever-increasing amount of cars in circulation [121]; and, they contribute to a healthy lifestyle [94]. Particularly throughout the COVID-19 pandemic, BSSs were considered a transportation alternative with a particularly low risk of user contamination [3, 92].

In this paper, we focus on dock-based BSSs, in which stations are located in different parts of the city, and from which commuters may rent and return bikes. While dock-based systems have several advantages (e.g. users get used to the location of stations and bikes), a main issue is that the station inventories may quickly become unbalanced, i.e., either rental demand cannot be met, given that not a sufficient number of bikes is available, or return demand cannot be met, when the station has no empty docks. Such imbalances often occur during rush hours on weekdays, when commuters relocate from their residential areas to the areas they work in the morning and do the return trip in the afternoon [82]. Unmet user demand likely causes user dissatisfaction, which the system operators seek to avoid as best as possible, given that it ultimately reduces the user base as the system's reputation is damaged.

An effective way to fight station inventory imbalances is to redistribute bikes among stations, a process known as *rebalancing*. The literature distinguishes two main types of rebalancing: user-based rebalancing and operator-based rebalancing. The former consists of incentives given to the users in order to return bikes at stations before they become empty [116]. In contrast, operator-based rebalancing is carried out by the BSS operators themselves, typically by dispatching trucks that relocate bikes between the stations. In this work, we focus on operator-based rebalancing, which has shown to be effective to increase demand satisfaction [see, e.g. 23] and is the common practice at major BSSs around the world, such as BIXI Montreal, Citi Bike in New York City and Ecobici in Mexico City. Such rebalancing is also a less expensive solution compared to installing more stations or adding more docks to already existing stations [109].

In most dock-based BSSs with operator-based rebalancing, the decision to actively rebalance

a station depends on which stations are considered *unbalanced*. Depending on the BSS, the criteria may be different for a station to be categorized as such. For example, at NiceRide (Minneapolis, U.S), a station is considered to be unbalanced when it is either completely empty or completely full [122]. The operators of Vélo'v (Lyon, France) classify a station as unbalanced if the absolute difference between the number of arrivals and departures is larger than the standard deviation of the distribution of these values over all the stations [5]. BIXI Montreal uses *inventory intervals* that establish an acceptable quantity of bikes at each station. Inventory intervals are manually set by BIXI's dispatching team, based on their experience with the station location, intraday demand fluctuation, and the day of the week.

Nonetheless, the rebalancing process itself remains costly, as it accounts for gas, the maintenance of the truck fleet, drivers' salaries, etc. In addition, it reduces the favorable impact that bike sharing claims to have on the environment. All considered, having a fleet of trucks large enough to rebalance all unbalanced stations at every hour is not financially viable for most BSSs, especially during peak hours. According to JCDecaux, a company that offers self-service bikes to different cities around the world, the estimated cost in 2009 to relocate a single bike within a BSSs was about three dollars [23]. A fleet of trucks available for rebalancing is therefore limited in size and cannot rebalance all unbalanced stations. It becomes imperative that the selection of stations to rebalance is carried out as effectively as possible.

Consider Figure 3.1, showing BIXI's (estimated on historical data) maximum hourly rebalancing capacity and the average daily number of unbalanced stations per hour for weekdays in July and August of 2019 (left) and 2020 (right). Throughout this period, the number of unbalanced stations consistently exceeded BIXI's maximum rebalancing capacity of approximately 46 stations per hour in 2019 and 22 stations per hour in 2020. As a consequence, the operator must select a subset of these stations to be rebalanced.

Ideally, the subset of unbalanced stations should be selected such that the number of served future demand requests is maximized, which requires an appropriate demand forecast. However, predicting the demand of a BSS is a complex task depending on several factors, such as the weather, the hour of the day, the day of the week, holidays and public events.

Demand prediction in BSS became particularly challenging in 2020 due to restrictive measures imposed by the governments in response to the COVID-19 pandemic, leading to a large part of the population working from home. Figure 3.2 shows the average number of trips and the average number of rebalancing operations as reported by BIXI during the weekdays in July and August 2019 (left) and 2020 (right). Not only did the number of trips in 2020 decreased considerably with respect to the same period in 2019, but the trip behavior also changed. In

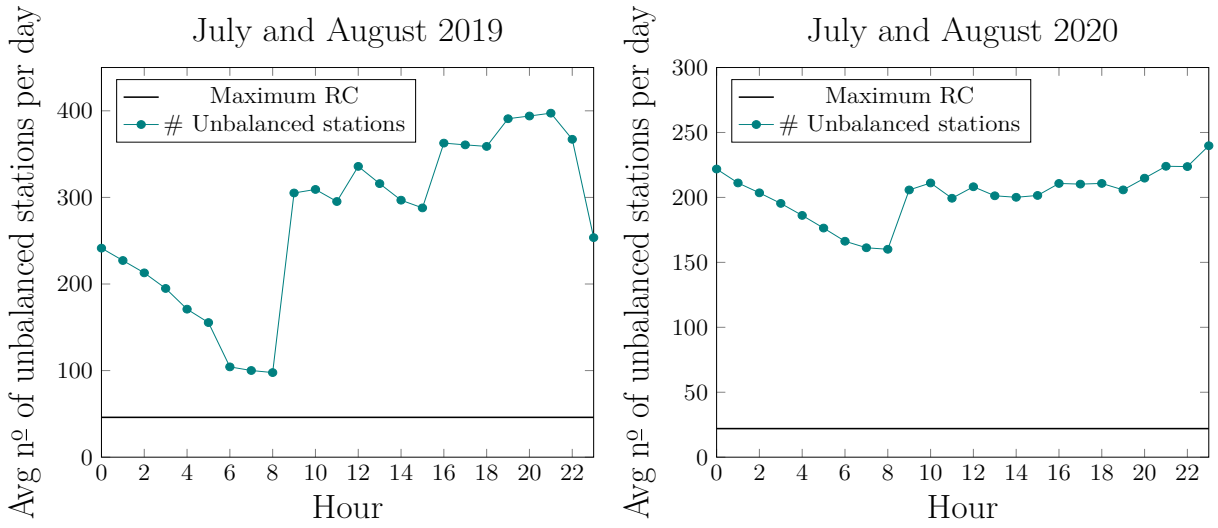

Figure 3.1 Maximum hourly rebalancing capacity and average daily number of unbalanced stations per hour at BIXI BSS during weekdays in July and August 2019 and 2020.

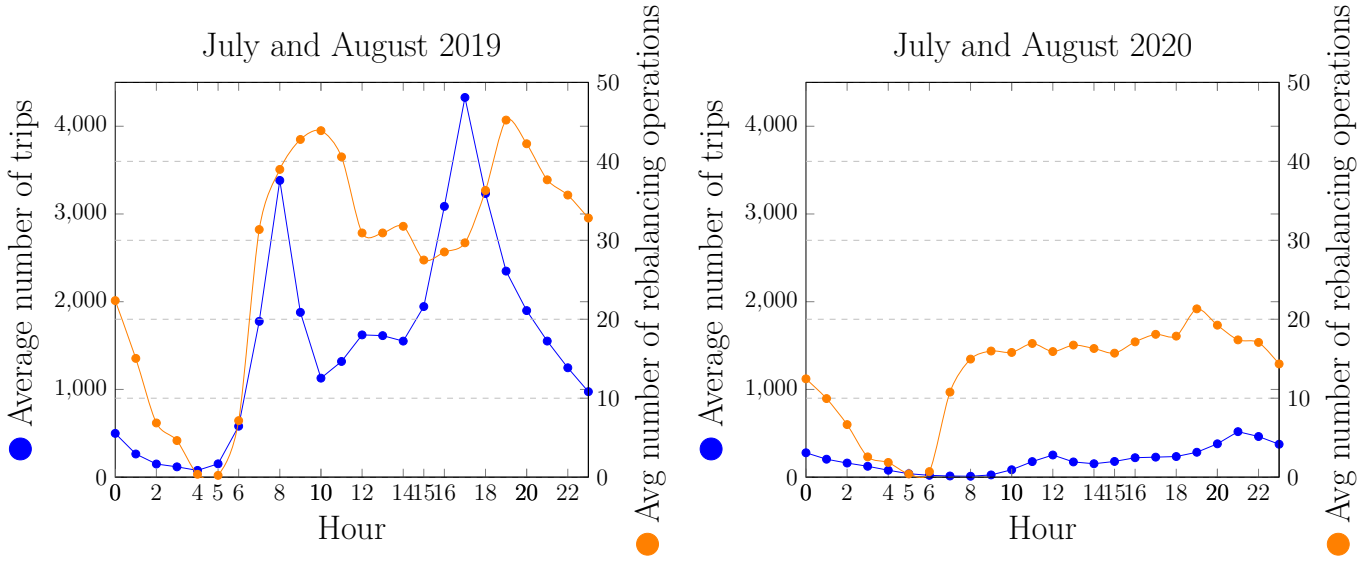

Figure 3.2 Average number of trips and rebalancing operations for all stations per weekday for July & August 2019 (left) and 2020 (right).

2019, the peak hours occurred right before and right after the working hours, i.e., at 8 a.m. and 5 p.m., respectively. However, this pattern was no longer observed in 2020, resulting in a flatter demand along the day.

Such drastic demand changes as observed from 2019 to 2020 severely affect the subset of stations that become unbalanced over time, complicating the demand forecast, and, in turn, affecting the choice of stations to be rebalanced. As a result, the pattern of BIXI's rebalancing operations changed drastically, reducing their activities by nearly half during the peak hours in 2020. Given that the manual planning is mostly based on the previous experience of the dispatching team, there is a significant risk that manually adjusted rebalancing strategies are ineffective in practice in a different environment of trip demand. This suggests that data-driven strategies that quickly adapt to changing demand hold certain benefits to assist BSS rebalancing operations.

In this paper, we focus on the data-driven selection of a subset of stations that should be considered for rebalancing. Such selection is based on demand prediction models that primarily focus on optimizing user satisfaction while simultaneously minimizing service costs. To this end, we first propose a model that automatically generates inventory intervals considering both rentals and returns forecast. Thereafter, we propose four strategies to prioritize the unbalanced stations according to their need for rebalancing. These strategies are based on the current inventory levels at the stations, the predicted demand for the next hours and, in the last strategy, the location of the stations. Each strategy aims to tackle a specific issue in order to improve the BSS performance. The first strategy focuses on prioritizing stations that are the most likely to generate lost demand. The second strategy aims to prioritize stations according to the amount of lost demand that can be avoided if they are immediately rebalanced. The third strategy prioritizes stations according to how unbalanced they are based on their inventory intervals. Finally, the fourth prioritization strategy further considers the proximity among unbalanced stations in order to favour the rebalancing of stations close to each other, leading to smaller operational routing costs for *a posteriori* truck route optimization. Throughout our prioritization strategies, we ensure that the set of stations selected for rebalancing results in a balanced number of bikes to be added and removed from them.

Given the lack of literature with a similar objective as ours (i.e., selecting a subset of stations), our four prioritization strategies are compared to a systematic approach, which emulates the prioritization strategy currently employed at BIXI. More specifically, the comparison is conducted by a tailored discrete-time simulation that computes the estimated lost demand (i.e., demand that could not be satisfied), the total number of alerts raised each time a station

becomes unbalanced, and the number of performed rebalancing operations in the system. Our computational experiments reveal that one of the proposed strategies is able to reduce the estimated lost demand by 35% for the 2019 season data, and by 65% for the 2020 season data, as compared to a baseline strategy that emulates the prioritization performed by the studied BSS. Besides, another prioritization strategy has shown effectiveness in selecting stations that are naturally grouped together while maintaining satisfactory performance measures. As outlined above, selecting a subset of critical stations that is aligned with the total rebalancing capacity is a primary task and concern of many system operators and, as such, can easily be integrated into the company's decision process. All proposed strategies are easy to implement and computationally cheap. They therefore provide an attractive alternative to more elaborate BSS planning approaches, such as those based on optimization models, which tend to be computationally challenging (or even intractable), given that they require to consider all stations within the same model.

The remainder of this paper is organized as follows. Section 4.2 reviews the most relevant literature in the area of rebalancing and prioritization strategies for bike-sharing systems. Section 4.3 describes how the inventory intervals are defined so as to serve as input to the prioritization strategies. Section 4.5 describes the different strategies proposed to score the rebalancing priorities of unbalanced stations. Section 3.5 presents and analyzes the computational experiments. Finally, Section 3.6 concludes the paper.

3.2 Literature Review

Rebalancing in BSSs can be divided into two main steps: (a) tactical inventory management, and (b) operational bike repositioning. Step (a) aims to establish the number of bikes in each station to meet the predicted demand as best as possible. Step (b) focuses on the actual dispatching operations that are necessary to achieve the desired inventory levels at the stations in order to rebalance the system.

In order to define, in step (a), the optimal inventories that are likely to provide sufficient bikes and free docks to satisfy future demand, it is necessary to forecast the latter sufficiently well. Trip demand is influenced by numerous external factors, such as the weather, the day of the week, the time of the day, land use, the location of the stations, points of interest, and social-demographic characteristics [25, 46]. Most of the proposed approaches to predict trip demand are either based on machine learning [see, e.g. 31, 53, 127] or on statistical models [see, e.g. 5, 15, 25, 39]. These models differ from each other in terms of the predicted time horizon (hourly, daily, or weekly), as well as on the geographic granularity of the predictions (station-level, cluster-level, or network-level). For instance, [127] and [39] predict the total

demand in the network for each observed hour. [31] and [5] predict the total demand for clusters of stations, while [15] estimate the probability that stations in a cluster become either completely full or empty. [25] propose a model that estimates the future demand for each station for five time periods along the day (morning, midday, afternoon, evening, and overnight). [53] predict the rentals and returns for each station per hour, using temporal (day, day of the week, holiday, etc.) and weather (temperature, humidity, rain, etc.) features. The authors also propose a reduction technique for the trip data, improving the computational execution time and erasing outliers from the dataset. For practical purposes, station-level demand predictions for shorter time-periods (such as one hour) seem preferable given that (i) the demand can drastically change from one hour to the next, and that (ii) the rebalancing process is actually planned and carried out at station-level.

Once the demand is properly predicted, optimal target inventory values can be determined. [106] model the station inventory by means of a Poisson queuing system that estimates its optimal number of bikes while ensuring a given service level. Also using a queuing system, [50] compute the target inventory values for the highest-demand stations, denoted central stations, considering their demand as well as those of their nearby stations. In the work of [72], inventories are optimized in order to maximize the amount of time a station is considered balanced. In [97], they estimate the optimum inventory by minimizing a user dissatisfaction function with penalty variables to control the weight given to the rental or return dissatisfaction. Likewise, [53] introduces a hyperparameter to prioritize either the rental or the return service level when computing the target inventory value of a station. [21] obtain target inventory values by minimizing the journey dissatisfaction, which is measured by the number of commuters who wait for missing bikes (for rentals) or docks (for returns), or by the number of times they change of station.

Regarding the operational decision-making step (b), the works in the literature can be categorized into two classes: those that assume that rebalancing operations are performed by the users of the system (under some incentive) and those that rebalance by means of a truck fleet coordinated by the operator. [14] propose a reward mechanism to encourage users to return bikes to certain stations in the BSS. In [33], the authors conclude that encouraging the users to return the bikes to a non-saturated station does not significantly improve the system's performance. However, they also show that the performance can be improved by constantly stimulating users to return bikes to a nearby station with lower inventory.

In the case of dock-based BSSs, as the one here considered, operator-based rebalancing via trucks has typically been modeled via mixed-integer linear programming [see, e.g. 2, 8, 10, 13, 20, 22, 26, 73, 88, 90]. These models generally aim to finding optimal truck routes

to rebalance a set of stations, typically seeking to maximize customer satisfaction. The latter may be achieved by minimizing the total lost demand [see, e.g. 2, 20, 73], keeping station inventories close to their respective target inventory values [see, e.g. 8, 90], or even by optimizing several (possibly conflicting) objectives [see, e.g. 86]. Unfortunately, the use of such models in practice is rather challenging, given that the resulting optimization models tend to be hard to solve. This typically limits their use to a small number of stations, given that intraday planning typically requires decisions within a matter of minutes. [106] observe that their formulation becomes difficult to solve even for small instances with 50 stations and 3 trucks. The authors, therefore, propose a heuristic that clusters stations using a maximum spanning star and then rebalances among clusters. Likewise, [43] cluster nearby stations and then rebalance among the clusters, where the capacity and the inventory of each cluster are given by the sum across its stations. While such an approach improves computational feasibility, it is based on the assumption that nearby stations have similar patterns and that rebalancing within each cluster is time feasible. Several other heuristic methods have been proposed [see, e.g. 77, 91, 99]. In particular, [117] propose a large neighborhood search algorithm that optimizes the truck routes only for stations that have raised an alert to the system. An interesting characteristic of their optimization model is that such alerts have different priorities which are proportional to their importance in the objective function. The list of stations to rebalance (i.e., those that raised an alert) along with their associated priorities have to be provided as input to the optimization model.

Our proposed prioritization strategies carry out step (a) and a preliminary step to facilitate step (b). While our approach does not directly optimize bike repositioning through the routing of dispatched trucks, we propose how to reduce the set of stations that should be considered in the network. By focusing on those that are the most urgent, we significantly facilitate efficient rebalancing operations. Indeed, a prior selection of stations can be very useful to scale optimization models to large BSSs by restricting the number of stations to be actually considered. Moreover, such approaches, based on alerts raised for stations that are susceptible to become unbalanced, are also easier to fit into existing practices at several BSSs, which often plan the dispatching operations based on such alerts. Nonetheless, because of limited resources in practice, planners are often required to choose a subset of the stations to rebalance.

The strategies proposed in the next sections seek to recommend the best subset of stations to be rebalanced over a prespecified period of time (e.g. one hour). Given that these prioritization strategies are easy to implement and computed within a matter of seconds, they provide an attractive alternative to computationally expensive optimization models and can easily complement systems that perform rebalancing planning based on raised station alerts.

We reiterate that to the best of our knowledge, there is a lack of existing literature that shares our precise objective. Note that clustering methods work by segmenting the entire network into smaller ones. Consequently, these methods assume that each subproblem is addressed independently by allocating the available trucks to specific zones. This approach indeed contrasts with our strategy of selecting a subset of unbalanced stations with the purpose of making the routing optimization problem computationally tractable while taking into account the entire network.

3.3 Inventory intervals

BSS operators (such as BIXI Montreal) often use intervals of acceptable inventory levels, referred to as inventory intervals. They are composed of a lower and an upper bound, as well as a target inventory value, which refers to the ideal inventory for that station and falls within the lower and upper bounds. Typically, each station has its specific inventory interval defined for a specific time period, and its values may change depending on the hour and day. When the inventory of a station falls outside of its specified inventory interval, the station is classified as unbalanced, which in turn triggers a rebalancing *alert*.

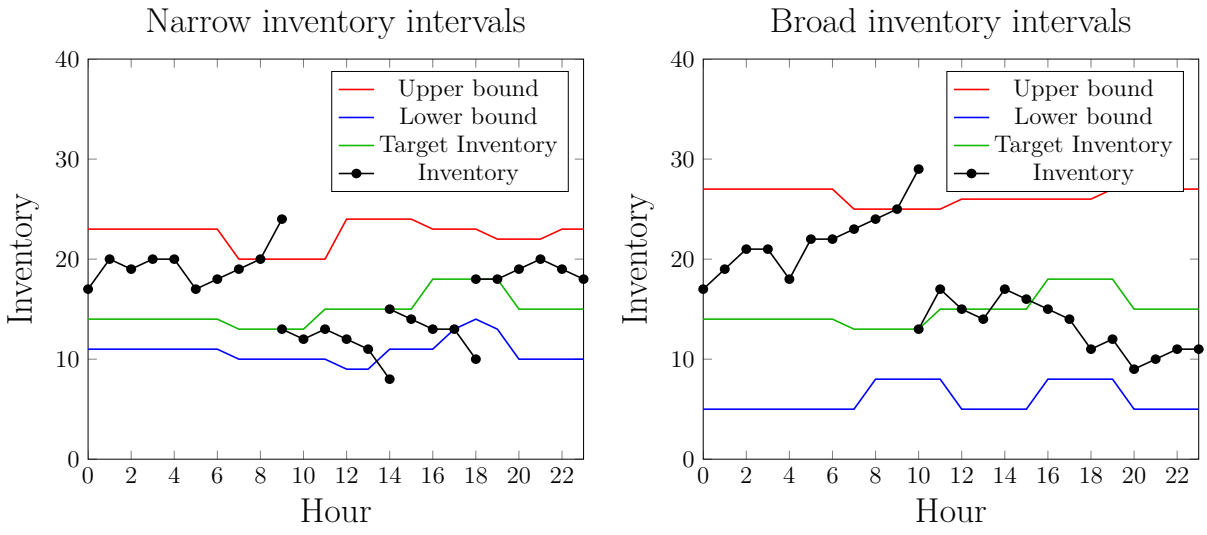

Figure 3.3 Trade-off between the number of alerts and lost demand when defining inventory intervals.

Figure 3.3 illustrates the evolution of the inventory of two stations with different inventory intervals and target values. Once the inventory is outside the interval, the station is assumed to be rebalanced such that, at the next hour, the inventory is set back to the target value. The width and the level of the inventory intervals have a major impact on the quantity

of lost demand, on the number of raised alerts, and ultimately, on the total number of rebalancing operations. For example, narrow intervals generate more alerts, but tend to keep the inventory closer to its target value, which in turn, decreases the likelihood of lost demand (see Figure 3.3, left). In contrast, wide inventory intervals create fewer alerts, but at the expense of keeping the inventory farther from its target values, which may increase the amount of lost demand (see Figure 3.3, right). Defining inventory intervals that make this delicate trade-off in order to maximize system performance over the entire day is therefore a sensitive challenge.

The model proposed here computes the inventory intervals and the target values based on the *service level* – which is defined as the proportion of satisfied trips at a given station [106]. In the literature, the metric used to estimate user satisfaction can be computed considering either rentals or returns individually [see, e.g. 50, 53, 106], or, as in our case, in a combined manner in which both rentals and returns are considered [see, e.g. 21, 97].

The inventory of the stations is typically modeled as an M/M/1/K queue, whose parameters are the time between the rentals, the time between the returns, the number of servers (here, a single station), and the maximum capacity of the server (i.e., the total number of docks C_s). As often assumed in the literature, we describe the trips by a Poisson distribution [40, 43, 53, 60, 98, 106, 109]. Hence, the times between rentals and returns follow an exponential distribution. To estimate the rentals and returns rate, we used the model presented by [53] which is able to predict the hourly trips at a station of the system by training a gradient-boosted tree (GBT) model and using a Singular Value Decomposition (SVD) method to reduce the dimension size and exclude outliers.

Given a station s with initial inventory f, a time period $[0, T]$, and the parameters mentioned above, we propose to compute the joint service level of s as:

$$SL_s(f,T) = \frac{\int_0^T \mu_s(t)(1 - p_s(f,0,t)) + \lambda_s(t)(1 - p_s(f,C_s,t))dt}{\int_0^T \mu_s(t) + \lambda_s(t)dt}, \tag{3.1}$$

where $p_s(f, N, t)$ is the probability that the station s has N bikes at time t, assuming that it had f bikes at time 0, and $\mu_s(t)$ (resp. $\lambda_s(t)$) is the expected demand at time t at station s for renting (resp. returning) bikes. The numerator of Equation (3.1) represents the expected number of satisfied trips at station s for the referred period and initial inventory value. After normalization, $SL_s(f,T) \in [0,1]$.

Thus, we can compute the minimum and maximum service levels for a station s in a time

period $[0, T]$ depending on the initial inventory at time 0 as follows:

$$SL_s^{min}(T) = \min_{f \in \{0,...,C_s\}} SL_s(f, T), \qquad \text{and} \qquad SL_s^{max}(T) = \max_{f \in \{0,...,C_s\}} SL_s(f, T). \tag{3.2}$$

Next, we establish a threshold Ω_s for the acceptable service level for a station s during time period $[0, T]$, defined as:

$$\Omega_s(T) = SL_s^{min}(T) + \beta(SL_s^{max}(T) - SL_s^{min}(T)), \tag{3.3}$$

where the hyperparameter $\beta \in [0, 1]$ controls how exigent the operator is about the network service. A small value of β approximates the threshold to the minimum service level, while a large β brings the threshold closer to the maximum service level. Here, $\Omega_s(T) \in [0, 1]$, given that each individual service level $SL_s \in [0, 1]$.

The inventory interval for station s for time period $[0, T]$ compatible with threshold $\Omega_s(T)$ is then defined as

$$\mathcal{I}_s(T) = \{f \in \{0, ..., C_s\} | \mathcal{L}_s(T) \leq f \leq \mathcal{U}_s(T)\}, \tag{3.4}$$

where

$$\mathcal{U}_s(T) = \max\{f \in \{0, ..., C_s\} | SL_s(f, T) \geq \Omega_s(T)\}, \text{ and} \tag{3.5}$$

$$\mathcal{L}_s(T) = \min\{f \in \{0, ..., C_s\} | SL_s(f, T) \geq \Omega_s(T)\}. \tag{3.6}$$

Finally, the target inventory value $\mathcal{T}_s(T)$ for station s for time period $[0, T]$ is given by the inventory value that provides the highest service level:

$$\mathcal{T}_s(T) = \arg\max_{f \in \{0,...,C_s\}} \{SL_s(f, T)\}. \tag{3.7}$$

Based on the above defined inventory intervals and the target inventory value, we next propose new strategies to effectively prioritize unbalanced stations that should be considered for rebalancing.

3.4 Prioritizing strategies for unbalanced stations

Operators typically follow a systematic approach to prioritize stations for rebalancing among those that have raised an alert. In particular, BIXI Montreal segments the unbalanced stations into three priority groups, each with specific criteria[2]. The first group, denoted

[2]This procedure was conceived after several exchanges with BIXI's planners. As such, it is not an official representation of BIXI's decision-making process.

K_1, contains stations classified as critical, i.e., stations that are either completely empty or completely full, and for which all their neighboring stations within a radius of 600 meters are also completely empty (or completely full)[3]. The second group, denoted K_2, includes unbalanced stations that are not part of K_1, but are located within a 600-meter radius of a metro station. This criterion recognizes the fact that stations located near metro stations tend to have a higher demand, since it is common for both modes of transportation to complement each other in the urban environment. The third group, K_3, comprises stations that are not part of the previous groups but are within 600 meters of any station in K_1 or K_2. Any unbalanced station not classified under K_1, K_2 or K_3 is not considered for rebalancing in the analyzed time period. The priority of stations in K_1 is always higher than that of stations in K_2, and the priority of stations in K_2 is higher than that of stations in K_3. Finally, within each group, stations are sorted based on their proximity to the nearest metro station. This means that stations closer to the metro have a greater likelihood of being selected for rebalancing.

In the following, we present our proposed prioritization strategies, that provide priority scores to the unbalanced stations according to different criteria.

3.4.1 Prioritization strategy based on inventory forecasting

The first prioritization strategy, denoted Pa_1, selects stations for rebalancing taking into consideration the current stations inventories, provided as input, and the expected demand, which is obtained by the gradient-boosted tree (GBT) regression model of [53]. We predict the inventory of each station s in the network at the next hour $t+1$ as follows:

$$\overline{f}_s^1(t+1) = f_s(t) + \mu_s(t) - \lambda_s(t), \tag{3.8}$$

where $\mu_s(t)$ and $\lambda_s(t)$ are, respectively, the expected demand for the number of rentals and returns at station s during time period $[t, t+1]$, and $f_s(t)$ refers to the actual number of bikes available at station s at the beginning of hour t. Note that (3.8) can also be extended to predict inventories for subsequent hours in advance. In this case, it suffices to predict bike demands for the hours of the considered time period.

Then, for each station s, the strategy computes two indices

$$\mathcal{B}_s^{1-}(t+1) = \max\{0, -\overline{f}_s(t+1)\} \qquad \text{and} \qquad \mathcal{B}_s^{1+}(t+1) = \max\{0, \overline{f}_s(t+1) - C_s\}, \tag{3.9}$$

[3]The radius of 600 meters is defined by BIXI based on the fact that an average person may walk this distance within 10-15 minutes, which is the time that BIXI considers acceptable for a commuter to walk seeking to be served.

which represent, respectively, the predicted inventory shortfall and surplus at hour $t+1$ at station s. By definition, only one of the two indices $\mathcal{B}_s^-(t)$ and $\mathcal{B}_s^+(t)$ can assume a value greater than 0 at the same time t. Finally, a prioritization score is computed for station s at hour t proportional to its predicted inventory shortfall or surplus, as follows:

$$\mathcal{P}_s^1(t) = \max\{\mathcal{B}_s^{1-}(t+1), \mathcal{B}_s^{1+}(t+1)\}. \tag{3.10}$$

As a result, strategy Pa_1 ranks and prioritizes stations in non-increasing order of their predicted inventory shortfall or surplus for the next hour(s).

3.4.2 Prioritization strategy based on inventory forecasting with immediate rebalancing

The second proposed strategy, denoted Pa_2, sorts unbalanced stations according to the amount of lost demand avoided by rebalancing operations. It assumes that a rebalancing operation at the beginning of hour t sets the inventory of a station to its target inventory value. Thus, the predicted inventory for each station s at the hour $t+1$ is computed as:

$$\overline{f}_s^2(t+1) = \mathcal{T}_s(t) + \mu_s(t) - \lambda_s(t), \tag{3.11}$$

The predicted shortfall and surplus of bikes at station s for the next hour are then computed as:

$$\mathcal{B}_s^{2-}(t+1) = \max\{0, -\overline{f}_s^2(t+1)\} \qquad \text{and} \qquad \mathcal{B}_s^{2+}(t+1) = \max\{0, \overline{f}_s^2(t+1) - C_s\}. \tag{3.12}$$

Finally, the prioritization score of station s at hour t is given by strategy Pa_2 as:

$$\mathcal{P}_s^2(t) = \mathcal{P}_s^1(t) - \max\{\mathcal{B}_s^{2-}(t+1), \mathcal{B}_s^{2+}(t+1)\}. \tag{3.13}$$

The first term of (3.13), i.e., $\mathcal{P}_s^1(t)$ represents the lost demand if station s is not rebalanced at hour t, whereas the second part, i.e., $\max\{\mathcal{B}_s^{2-}(t+1), \mathcal{B}_s^{2+}(t+1)\}$ refers to the lost demand if that station has its inventory set to its target value $\mathcal{T}_s(t)$ at the beginning of hour t. Note that, theoretically, $\mathcal{P}_s^2(t)$ may be negative, i.e., rebalancing a station may increase the lost demand. Throughout our computational experiments with real-world data, however, we have never observed negative values for $\mathcal{P}_s^2(t)$. Finally, the strategy outputs the stations in non-increasing order of their prioritization scores.

3.4.3 Prioritization strategy using inventory intervals

The third prioritization strategy we propose, denoted Pa_3, is more conservative than the strategies above. Instead of giving a high prioritization score to stations whose inventories are predicted below zero or above their dock capacity (i.e., yielding lost demand), Pa_3 prioritizes unbalanced stations according to the predicted deviation of the station's inventory from its inventory interval bounds (eqs. (3.5-3.6)). Thus, stations might be rebalanced before they actually start to generate lost demand (i.e., before they are completely full or empty). Strategy Pa_3 computes two indices:

$$\mathcal{B}_s^{3-}(t+1) = \max\{0, \mathcal{L}_s(t) - \overline{f}_s^1(t+1)\} \quad \text{and} \quad \mathcal{B}_s^{3+}(t+1) = \max\{0, \overline{f}_s^1(t+1) - \mathcal{U}_s(t)\}, \quad (3.14)$$

where $\mathcal{L}_s(t)$ and $\mathcal{U}_s(t)$ correspond, respectively, to the lower and upper bounds of the inventory interval of station s at hour t. We also recall that $\overline{f}_s^1(t+1)$ corresponds to the predicted station inventory at hour $t+1$ if no rebalancing is carried out at station s (see Section 3.4.1). The priority score for strategy Pa_3 is then computed as:

$$\mathcal{P}_s^3(t) = \max\{\mathcal{B}_s^{3-}(t+1), \mathcal{B}_s^{3+}(t+1)\}. \quad (3.15)$$

As a result, Pa_3 prioritizes stations that are expected to remain unbalanced in the next hour if the operator does not rebalance them. These stations are then sorted and prioritized in non-increasing order of the deviation between their predicted inventories and their inventory interval bounds.

3.4.4 Prioritization strategy based on neighbourhood

The fourth prioritization strategy proposed here, denoted Pa_4, diverts from the previous ones in the sense that it additionally takes into consideration geographical information to rank unbalanced stations. By favoring the rebalancing of neighbouring stations, Pa_4 is expected to obtain more compact and thus less costly (and faster) dispatching routes that group together nearby stations.

Let us define $\mathcal{N}_s$ as the index set of all stations within a radius of 600 meters from station s. The priority score computed by Pa_4 is given by:

$$\mathcal{P}_s^4(t) = \begin{cases} \gamma \times \mathcal{P}_s^*(t) + (1-\gamma) \times \frac{\sum_{s' \in \mathcal{N}_s, s' \neq s} \mathcal{P}_{s'}^*(t)}{|\mathcal{N}_s|}, & \text{if } \mathcal{P}_s^*(t) > 0 \\ \mathcal{P}_s^*(t), & \text{otherwise.} \end{cases} \quad (3.16)$$

Here, the priority score $\mathcal{P}_s^*(t)$ refers to any of the three priority scores $\mathcal{P}_s^1(t), \mathcal{P}_s^2(t)$ or $\mathcal{P}_s^3(t)$ presented above. The hyperparameter $\gamma \in [0, 1]$ controls the weight of the neighbourhood inventory information on the computation of the priority score. By using $\gamma < 1$, the priority score computed by Pa_4 for a station s also takes into consideration the priority scores of its neighboring stations. As γ approaches 0, those scores tend to prevail in (3.16), prioritizing stations surrounded by others that raised an alert. As before, the stations are then sorted in non-increasing order of their priority scores.

For all of the proposed prioritization strategies, stations with a priority score of 0 or lower are simply discarded. Among stations with the same score, priority is given to those closer to a metro station.

3.4.5 Post-processing to ensure a balanced selection of pick-ups and drop-offs

Our prioritization strategies are responsible for sorting the stations according to their rebalancing priority, but they do not take into consideration practical constraints such as the hourly rebalancing capacity of the operator or the parity between the amount of bikes added and removed from the BSS. Indeed, the latter is an important criterion for the rebalancing team in order to ensure that the number of bikes picked up is approximately the same as the number of bikes dropped off.

In view of that, we developed a post-processing procedure that follows the application of any of the previous prioritization strategies. This procedure has two steps and it is illustrated in Figure 3.4. In the first step, the procedure receives the list of sorted unbalanced stations and splits it into two sorted sub-lists of stations, namely: (i) a list $\mathcal{S}^{add}$ of stations at which bikes are added (i.e., dropped off), and (ii) a list $\mathcal{S}^{rem}$ of stations from which bikes should be removed (i.e., picked up). The amount of bikes added or removed from the stations is exemplified in the figure and corresponds to the number of required bikes to restore the inventory of the stations to their respective target values. In the second step, the stations are selected for rebalancing according to the value of the *accumulator* variable, which is initialized with 0. If $accumulator \geq 0$, the procedure selects the highest priority station from $\mathcal{S}^{rem}$ to be rebalanced, and updates the accumulator accordingly. Otherwise, the procedure selects for rebalancing the highest priority station from $\mathcal{S}^{add}$, and the accumulator variable is increased. Each time a station is selected in step 2, it is removed from its corresponding sub-list. The post-processing procedure is halted when either $\mathcal{S}^{add}$ or $\mathcal{S}^{rem}$ are empty, or when the number of rebalancing operations equals the BSS rebalancing capacity (i.e., the maximum number of rebalancing operations the operator is likely to be able to carry out). At the end, the procedure outputs the set of stations to be rebalanced in the simulation.

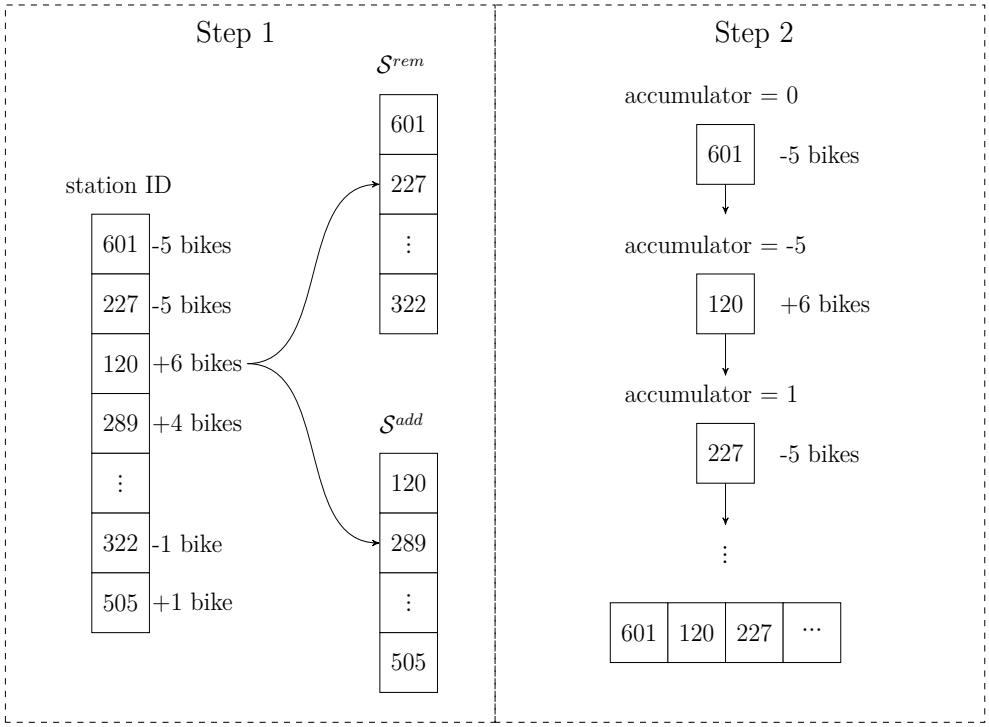

Figure 3.4 Illustration of the post-processing procedure is divided into two steps. Note that the original sorting of the stations provided by the prioritization strategy is preserved in the sub-lists.

3.5 Computational experiments

We now report on computational experiments that have been carried out to evaluate the performance of the proposed prioritization strategies. We first present the details of the dataset used in our study. In the sequel, we present the simulation used to estimate the inventory response to each of the applied rebalancing strategies. Specifically, it computes three performance metrics: lost demand, the number of raised alerts, and the number of performed rebalancing operations. These metrics are then compared for the different prioritization strategies. Finally, in a specific subsection, we assess Pa_4 with respect to the clustering property of its prioritized stations.

3.5.1 Data set

We consider real-world data from BIXI Montreal for the 2019 and 2020 seasons. The dataset used in the experiments contains hourly information for time, weather, trips, and stations. Time features hold temporal data such as the hour, day, whether it is a holiday, and the day of the week. Weather features store information such as temperature, wind speed and

relative humidity. Trip features are composed of the number of bike rentals and returns, observed at each station of the network. Finally, station features contain information regarding the geographical location of each station, as well as their corresponding number of docks. Time and weather features were collected from `https://climate.weather.gc.ca` (except for the holiday feature that was manually imputed), while trip and station features were both provided by BIXI Montreal. More details about the importance of the different features for the GBT used here can be found in [52].

The here considered BSS network from BIXI Montreal contains a total of 620 stations in 2019 and 641 stations in 2020. Each dataset was split into training, validation, and test data, and they each contain all stations from the corresponding year. The training and validation data are used, respectively, to fit the machine learning model parameters and tune its hyperparameters. Finally, the test dataset was used to provide an unbiased performance evaluation of the different prioritization strategies.

Because of a large observed discrepancy in the frequency and the behavior of trips in 2020, caused by the COVID-19 pandemic, we used different strategies for selecting the dataset for each simulated season. For the 2019 tests, the training dataset contains data from April 2018 to June 2019, minus the months during which BIXI is out of service (i.e., December, January, February, and March). The validation dataset is composed of data from the first 15 days of July and August 2019, and the test dataset uses data from the remaining days of July and August 2019. We opted to divide both July and August into validation and test datasets so that the model is less sensitive to demand changes observed between consecutive months.

Note that the physical network of BSSs typically changes over time, which makes it often difficult to use data linked to a specific station ID over longer periods of time. Therefore, we do not use data prior to the 2018 season. We have also observed that, for the 2020 season, adding training data from previous years deteriorates traffic prediction. Consequently, we used training data from April 2020 to June 2020. In that case, the validation dataset contains the first 15 days of July and August 2020, while the test dataset contains the remaining days of July and August 2020. Note that we focus our experiments on the months of July and August because of their high demand and importance throughout the season.

3.5.2 Simulation

Our simulation aims to emulate inventory fluctuations in the BSS according to the trip demand and the performed rebalancing operations. Thus, we can evaluate the proposed prioritization rebalancing strategies regarding the measured performance metrics.

The simulation starts by initializing the inventory of all stations with their respective target values. By doing so, we guarantee that the experiment starts with an optimal number of bikes in the system. Then, at every simulated hour, the inventory of all stations is modified according to the historical rentals and returns. We assume that all trips at each hour take place simultaneously to stress the system, thus capturing its possible failures. In the sequel, the simulation launches the rebalancing process by gathering the stations that raised alerts. The set of unbalanced stations is then forwarded for prioritization and, after being sorted, to the post-processing procedure, which ensures that the total number of bike pick-ups and bike drop-offs is sufficiently close to each other. The latter then returns a subset of the stations to be rebalanced. The number of stations within this subset is limited by the total rebalancing capacity, which is given as input to the simulator and represents the total number of stations the operator is capable of rebalancing at each hour. After the prioritized stations have been rebalanced by setting their inventories to their target values, the simulation proceeds to the next simulated hour until all hours are iterated.

After execution, the simulator returns three metrics that are important to the BSS operator: (i) the number of raised alerts, indicating how often a station has been classified as unbalanced, thus representing how stressed the system was; (ii) the number of rebalancing operations, allowing for an estimation of the operational rebalancing costs; and (iii) the amount of rental and return requests that could not be satisfied, directly affecting the customer satisfaction.

The pseudo-code of our simulator, as well as its execution pipeline and description can be found at `github.com/clara91/Data-driven-prioritization-strategies-for-BSSs`.

3.5.3 Simulation results

We now report on the simulation results of BIXI's data for the 2019 and 2020 seasons. We here consider that the maximum rebalancing capacity per hour was 46 for the 2019 season and 22 for the 2020 season. These numbers correspond to the average number of rebalancing operations performed by BIXI during the peak rebalancing hours in the two observed years.

All prioritization strategies have been fed with the same inventory intervals in order to ensure a fair comparison of their performance. Even though the inventory intervals used by BIXI were available to us, incorporating them into the analysis may have compromised the validity of our conclusions.

As shown in Section 4.3, the value of parameter β plays an important role in the definition of the inventory intervals and directly impacts the performance measures. Large values of

β result in narrow inventory intervals, and consequently, in a larger number of alerts, while potentially improving demand satisfaction. In contrast, small β values lead to wide intervals, decreasing the number of raised alerts at the stations, but possibly decreasing demand satisfaction. We, therefore, evaluated the prioritization strategies using three values for β, specifically $\beta = 0.75$ (narrow intervals), 0.50 (medium intervals), and 0.25 (wide intervals).

Evaluation of Pa_1, Pa_2 and Pa_3

We first focus on the comparison of the first three prioritization strategies Pa_1, Pa_2 and Pa_3, as well as the emulated strategy employed by BIXI Montreal. Table 4.2 provides an overview of the total lost demand (expressed as a percentage of the total demand, i.e., rental plus return demand) calculated across the entire network for the 768 simulated hours from the test dataset. The table also includes the hourly average of raised alerts and the hourly average number of rebalancing operations performed in the system over the same 768 simulated hours. The presented results were collected from the simulation using the proposed prioritization strategies Pa_1, Pa_2, and Pa_3, as well as the emulation of BIXI's prioritization strategy. Finally, for each of the presented metrics, we report the relative difference (column $\Delta(\%)$) calculated with respect to the values obtained using BIXI's prioritization strategy.

Table 4.2 suggests the following conclusions:

1. **General impact on lost demand.** The three proposed strategies consistently resulted in lower levels of lost demand compared to BIXI's prioritization strategy across all simulated scenarios. This suggests that the proposed strategies excel at prioritizing stations for rebalancing.

2. **Best strategy to reduce lost demand.** Among the proposed prioritization strategies, Pa_3 is the most effective strategy to reduce lost demand and the number of raised alerts, achieving reductions of $\approx 35\%$ in lost demand for 2019, and $\approx 65\%$ in 2020. However, it's important to consider the tradeoff involved when using Pa_3 as it requires a higher number of rebalancing operations compared to BIXI's prioritization strategy. Specifically, even though the performed rebalancing respects the maximum rebalancing capacity, more rebalancing operations also translate into higher operational costs.

3. **Best strategies with moderate rebalancing.** Pa_1 and Pa_2 reduce the lost demand without necessarily increasing the number of rebalancing operations. In fact, they often led to fewer rebalancing operations in the system – for all β values in 2019, and 2 out of 3 in 2020. This is explained by the fact that both Pa_1 and Pa_2 assign a positive score to unbalanced stations only if they are predicted to generate lost demand in the

Table 3.1 Performance metrics for the 2019 and 2020 seasons.

Season	Prioritization Strategy	β	Lost Demand		Alerts		Rebalancing	
			total %	Δ(%)	per hour	Δ(%)	per hour	Δ(%)
2019	BIXI	0.25	5.22		69.86		23.86	
		0.5	4.96		95.64		29.48	
		0.75	5.13		142.41		35.87	
	Pa_1	0.25	4.19	▼ -19.72	69.34	▼ -0.74	23.35	▼ -2.14
		0.5	4.05	▼ -18.23	109.83	▲ 14.84	23.83	▼ -19.17
		0.75	4.00	▼ -22.03	170.38	▲ 19.64	24.02	▼ -33.04
	Pa_2	0.25	4.16	▼ -20.24	69.03	▼ -1.19	23.20	▼ -2.77
		0.5	4.05	▼ -18.35	110.01	▲ 15.03	23.63	▼ -19.84
		0.75	3.95	▼ -22.95	169.87	▲ 19.28	23.89	▼ -33.40
	Pa_3	0.25	3.99	▼ -23.46	53.54	▼ -23.36	27.19	▲ 13.96
		0.5	3.58	▼ -27.71	74.80	▼ -21.79	31.51	▲ 6.89
		0.75	3.33	▼ -35.13	116.63	▼ -18.10	36.03	▲ 0.45
2020	BIXI	0.25	2.47		43.60		9.55	
		0.5	2.13		57.22		11.42	
		0.75	2.10		88.21		14.59	
	Pa_1	0.25	1.57	▼ -36.44	31.74	▼ -27.20	10.93	▲ 14.45
		0.5	1.53	▼ -28.17	51.98	▼ -9.16	10.97	▼ -3.94
		0.75	1.53	▼ -27.14	92.29	▲ 4.63	10.99	▼ -24.67
	Pa_2	0.25	1.55	▼ -37.25	32.02	▼ -26.56	10.92	▲ 14.35
		0.5	1.53	▼ -28.17	51.98	▼ -9.16	10.97	▼ -3.94
		0.75	1.52	▼ -27.62	92.65	▲ 5.03	11.06	▼ -24.19
	Pa_3	0.25	1.15	▼ -53.44	20.64	▼ -52.66	12.53	▲ 31.20
		0.5	0.88	▼ -58.69	28.26	▼ -50.61	14.07	▲ 23.20
		0.75	0.72	▼ -65.71	48.87	▼ -44.60	16.47	▲ 12.89

next hour, while Pa_3 bases its prioritization score on the violation of their inventory intervals. As a result, Pa_3 prioritizes a larger number of stations with scores greater than 0, leading to an increased frequency of rebalancing operations compared to Pa_1 or Pa_2.

4. **Impact of inventory interval size.** For our prioritization strategies, the relative reduction of the lost demand was more substantial when narrow inventory intervals ($\beta = 0.75$) were employed, generally leading to a higher number of alerted stations and slightly more rebalancing operations. In contrast, this behavior has not been observed during the 2019 season for BIXI's prioritization strategy. Here, narrow inventory intervals ($\beta = 0.75$) resulted in more alerts and rebalancing operations, but this did not translate into a lower lost demand. We analyze this curious behaviour below in Section 3.5.3.

5. **Pattern shift from 2019 to 2020.** As a result of the lockdown measures applied by the Canadian authorities in response to the COVID-19 pandemic, user behavior drastically changed. Most likely attributable due to the increased work-from-home, the demand peaks shifted and the total number of trips reduced by 85% (e.g., from about 4000 trips to about 600 trips at the respective peak hours). As a result, the relative lost demand also dropped by at least 50% from 2019 to 2020, as stations were substantially less stressed and users were less often faced with empty or full stations.

6. **Benefits of data-driven methods.** Even though the relative lost demand reduced from 2019 to 2020, this was expected (see previous point) and it is questionable whether the BIXI's rebalancing strategy has adapted to the new demand pattern in an ideal manner. Our prioritization strategies demonstrated a substantially higher reduction of the relative lost demand over BIXI's strategy in 2020. Such an improvement, despite the fact the total demand was much lower, is remarkable. It is, in fact, much more difficult to further reduce lost demand when the total number of trips is low. For example, 5% of lost demand at a peak-hour in 2020 may refer to about 15 out of 600 trips, while it refers to 200 out of 4000 trips in 2019. Reducing an abundant lost demand is easier than a sparse one, since the latter requires a very precise identification of the stations at which lost demand can be further reduced. It is therefore even more impressive to see the lost demand in 2019 reducing to about 0.72% for strategy Pa_3. This does not only highlight that our prioritization strategies are effective, it illustrates their ability to adapt to new demand patterns and, as such, highlights the importance of data-driven strategies in general that can adapt to changes in demand patterns much faster than manual adjustments .

Trade-off between the number of alerts and the rebalancing capacity

BIXI's prioritization strategy presenting higher lost demand with tighter inventory intervals (β = 0.75) is unexpected, given that our prioritization strategies consistently reduced lost demand as the inventory intervals got tighter. It turns out that the trade-off between the number of alerts and the rebalancing capacity is an important one.

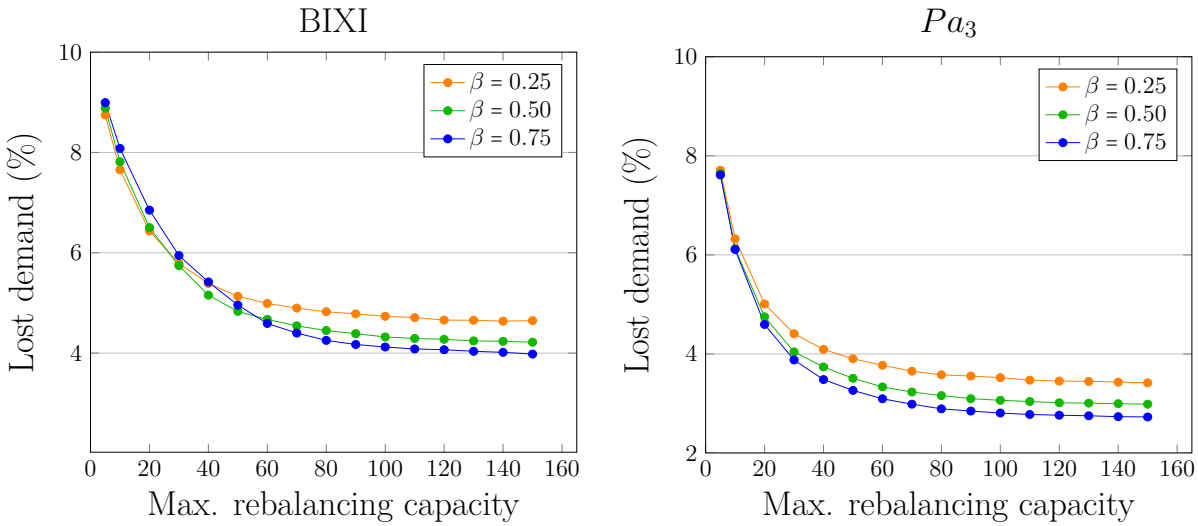

Figure 3.5 Percentage of lost demand as a function of the rebalancing capacity of the system on the 2019 test dataset.

Figure 3.5 presents the lost demand as a function of the BSS rebalancing capacity for both BIXI's prioritization strategy (left) and Pa_3 (right) throughout the 2019 test dataset. For strategy Pa_3, more narrow intervals (i.e., higher β values) consistenly lead to lower lost demand. In contrast, for BIXI's prioritization strategy, narrow intervals only perform well when the rebalancing capacity is sufficiently high (more than 50) to deal with the large amount of raised alerts. When stations are stressed (i.e., there is a high number of raised alerts), BIXI's prioritization strategy has difficulties identifying the most critical stations, which leads to higher lost demand when the rebalancing capacity is low (here, less than 40). Here, wider inventory intervals raise alerts only for the most unbalanced stations, which is more aligned with the limited rebalancing capacity. Indeed, the rebalancing capacity used in our case-study in Table 4.2 assumes a rebalancing capacity of 46, which lies at the threshold where more narrow intervals are beneficial when using BIXI's prioritization strategy.

This analysis allows for deriving two key insights to BSS operators, particular those with a limited rebalancing capacity (which is an economical key concern for operators). First, it highlights the importance of an effective prioritization strategy, particularly when the

number of alerts is high. Prioritizing the wrong stations is as ineffective as raising the wrong alerts. Second, being capable of adjusting the number of alerts (by smartly selecting the intervals and choosing an appropriate interval size) is crucial, since the operator may want to keep the number of alerts aligned with the rebalancing capacity, particularly when the subsequent prioritization process is not robust to changes in the demand pattern. The prioritization strategies here presented improve over BIXI's prioritization strategies in both concerns, which makes them an attractive tool to integrate while working towards more effective and automated rebalancing strategies.

Evaluation of Pa_4

Recall that all proposed prioritization strategies (i.e., Pa_1–Pa_4) respect the total rebalancing capacity prescribed by the operator (i.e., the maximum number of stations that can be selected for rebalancing). Further, we ensure in the post-processing procedure (see Section 3.4.5) that the number of bikes to be picked up and dropped off at the selected stations is approximately the same. While this may facilitate the planning of feasible rebalancing routes of the vehicles, stations may still be quite far from each other. Selecting a subset of stations that is naturally more clustered may therefore facilitate the generation of vehicle routes that rebalance the selected stations efficiently.

To this end, in this section, we turn our analysis to the prioritization strategy Pa_4, which is designed to favour the rebalancing of proximal unbalanced stations as γ tends to zero. Thus, it is expected that less lengthy, and consequently, less expensive routes can be derived for the trucks that are in charge of the system rebalancing. We report in this section results of strategy Pa_4 with $\mathcal{P}_s^3(t)$ in eq. (3.16). The results of Pa_4 obtained with the other prioritization algorithms are similar to the ones reported here in this section.

We first introduce an additional metric to measure the compactness of a set of stations prioritized for rebalancing. Let $G = (N, E)$ be a graph for which there is a node $n_i \in N$ corresponding to each station selected for rebalancing at a given hour. Set E contains all edges e_{ij} such that $n_i \in N$ and $n_j \in N$, and the distance between the stations represented by n_i and n_j is not more than 600m. Let us also denote $S_i \subseteq N$ as the set of nodes connected to $n_i \in N$. We then compute the *Watts–Strogatz clustering coefficient* [124] of G as:

$$\eta_G = \frac{1}{|N|} \sum_{n_i \in N} \eta_i, \quad \text{where } \eta_i \text{ is computed as} \quad \eta_i = \frac{\left|\{e_{jk} \in E : n_j \in S_i, n_k \in S_i\}\right|}{\binom{|S_i|}{2}}. \tag{3.17}$$

The Watts–Strogatz clustering coefficient measures the inherent tendency of a graph to form clusters. In fact, the value η_i of a node $n_i \in N$ measures how close its neighbours are to

forming a clique.

Figure 3.6 illustrates four graphs constructed from the stations prioritized by Pa_4 in our simulation run for July 19, 2019 at 11 am, using $\beta = 0.75$ and $\gamma = 0.25, 0.5, 0.75$ and 1. Nodes in N are drawn as red dots, whereas edges in E are indicated as yellow lines. We note from the figure that small γ values yield graphs for which the selected stations can be more naturally clustered, yielding higher Watts–Strogatz clustering coefficient values.

Next, Figure 3.7 illustrates the impact of the hyperparameter γ on the Watts-Strogatz clustering coefficient, as well as on the three performance measures used in our study. Moreover, we include in Figure 3.7 the best metric values obtained by the BIXI's prioritization strategy among those obtained with $\beta = 0.25, 0.5, 0.75$, which can be found in Table 4.2.

In particular, for each value of $\gamma \in \{0, 0.05, 0.10 \ldots, 1\}$, Figure 3.7a presents the average Watts-Strogatz clustering coefficient of the graphs built from the stations selected by Pa_4 around the demand peak hours of the 2019 BIXI's season, i.e., between 7-11 am and 4-8 pm. We can notice that η_G reaches its maximum for $\gamma = 0.15$, and not at $\gamma = 0$. Since the stations available for prioritization might vary over time due to previous rebalancing operations, being too greedy towards the prioritization of neighbouring stations might cause an absence of clusters of unbalanced stations for subsequent simulated hours. For reference, Pa_4 yields an η_G superior to that obtained by means of the BIXI's prioritization strategy for $\gamma \leq 0.5$, regardless of the tested value of β. Moreover, as expected, we can observe in Figures 3.7b-3.7d that Pa_4 improves its performance with respect to the other three metrics as γ approaches 1, as the influence of the station's neighborhood in the prioritization procedure is decreased. Note that Pa_4 is equivalent to Pa_3 when $\gamma = 1$.

Although the number of rebalancing operations observed with the use of Pa_4 is larger for smaller γ, such operations are potentially less costly since the selected stations are closer to each other. Such tradeoff between the clustering coefficient and the cost of the system's rebalancing operations indeed merits further investigation. However, it must be conducted in conjunction with routing optimization algorithms, which are out of the scope of the present paper. We leave this to future research.

3.6 Concluding remarks

In this paper, we proposed a series of strategies to select the unbalanced stations that should be prioritized for rebalancing in dock-based BSSs. This is an important issue for BSS operators, given that it is neither economically feasible nor ecologically desirable to provide a sufficiently large truck fleet in order to rebalance all unbalanced stations at the same time.

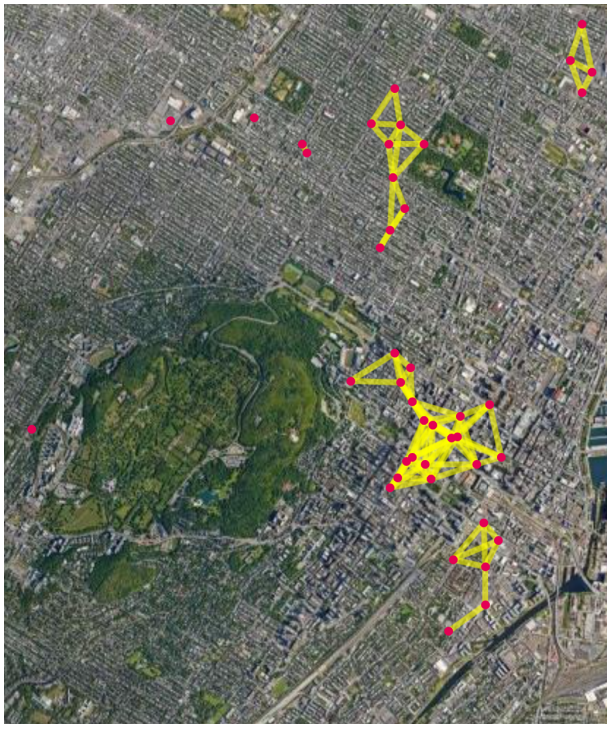

(a) $\gamma = 0.25, \eta_G = 0.59$

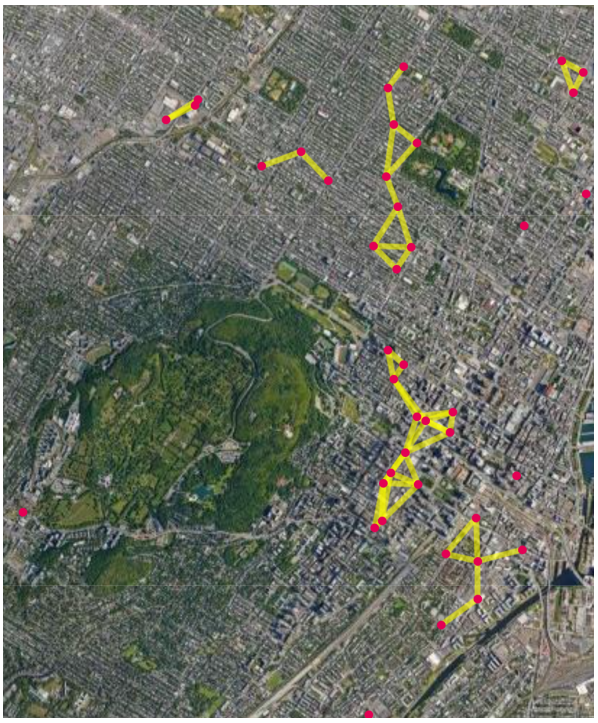

(b) $\gamma = 0.5, \eta_G = 0.48$

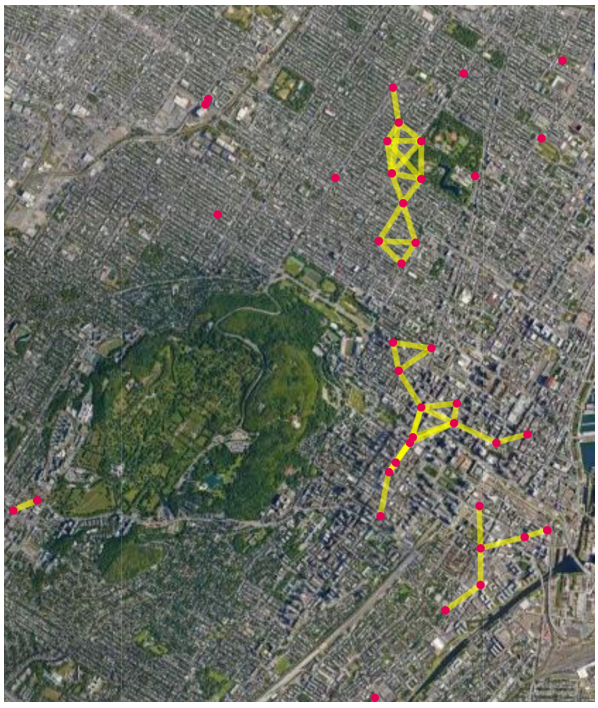

(c) $\gamma = 0.75, \eta_G = 0.29$

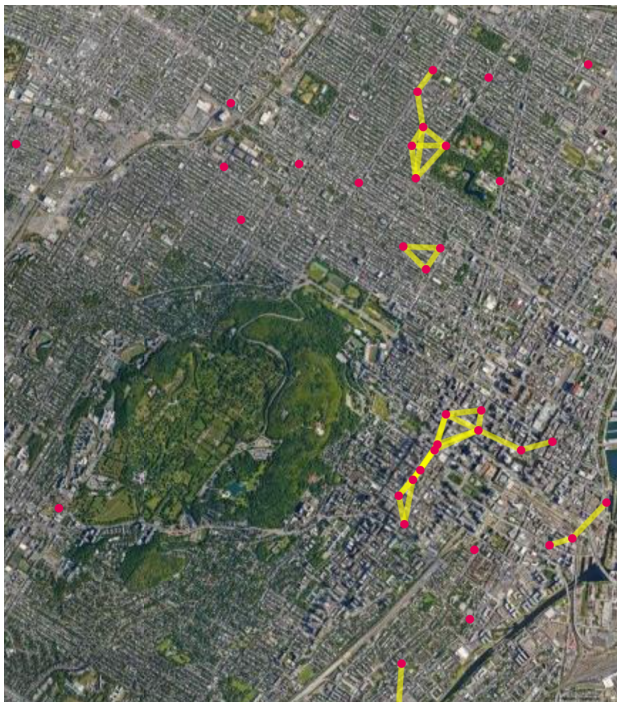

(d) $\gamma = 1, \eta_G = 0.27$

Figure 3.6 Stations prioritized by Pa_4 on July 19, 2019 at 11 a.m. with different values of γ.

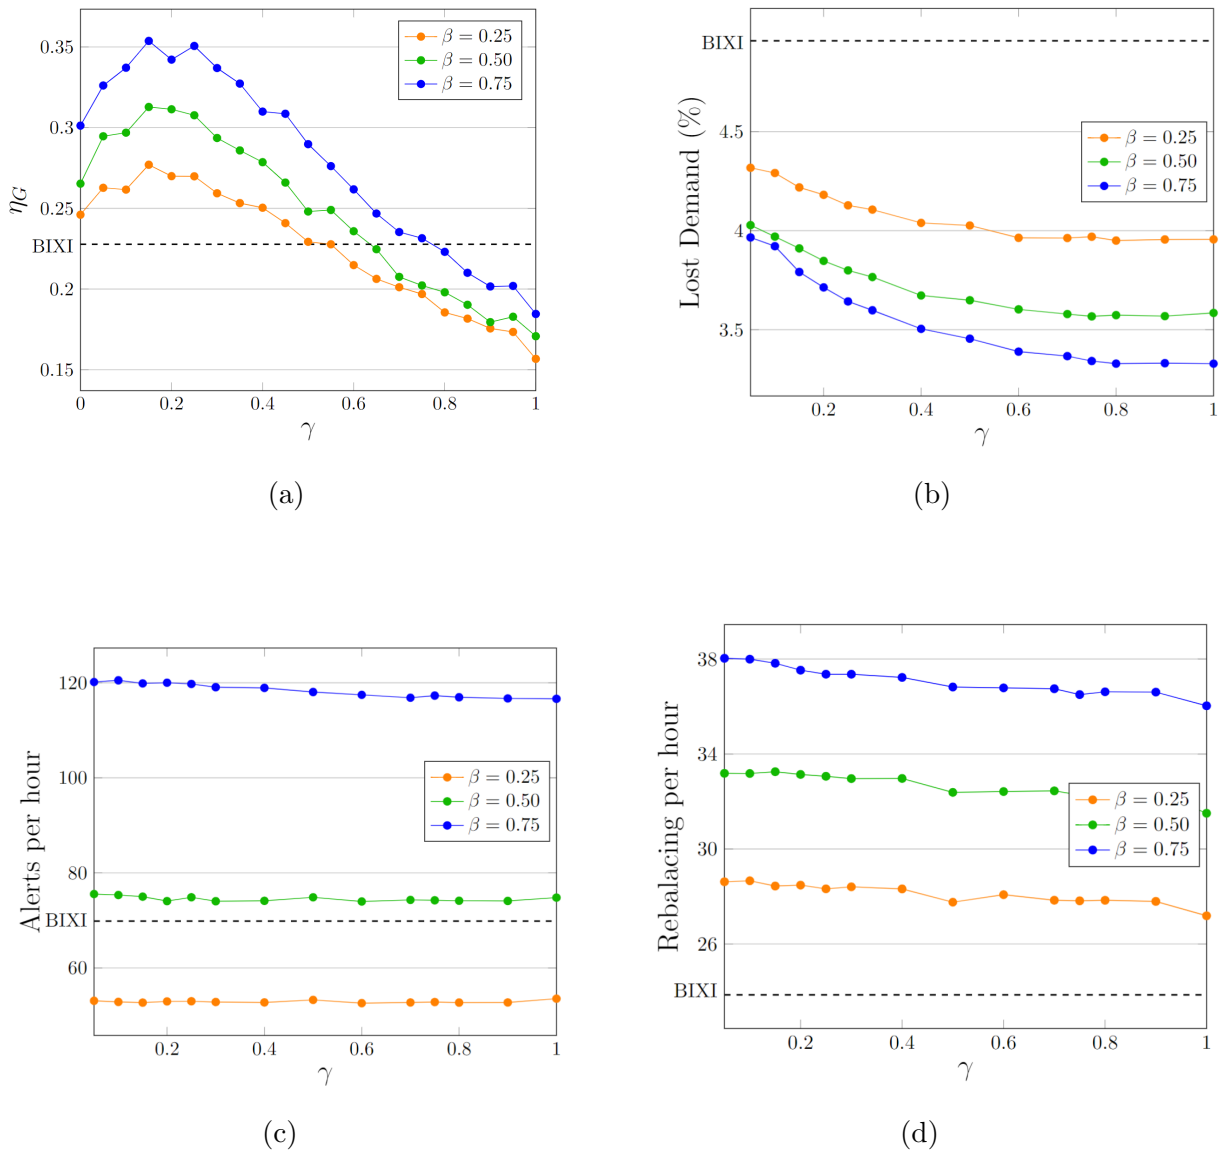

Figure 3.7 Pa_4 performance measures.

As such, we aim at improving a step that is present in many BSS operating processes, yet little addressed by the existing literature.

Our prioritization strategies make use of a wide range of indicators to select stations for rebalancing, including the forecast of future station inventories, the impact of potential rebalancing on station inventory, and the predicted deviation of the station inventories from their inventory intervals (the latter being automatically computed using demand prediction based on historical trips, weather, and temporal data). As inventory intervals are already currently used in practice by BSS operators, the implementation of our prioritization strategies into the practical information- and decision process is straightforward. Once the predictive model is trained, inventory intervals and prioritization strategies are computed in a matter of seconds, which is a practical requirement in the ongoing effort to rebalance stations throughout the day.

By comparing the proposed prioritization strategies, we observe that all of them are capable of reducing the occurrences of lost demand. This shows that a judicious choice of stations to be rebalanced has a considerable impact on the service level and, hence, customer satisfaction. Among the proposed strategies, Pa_1 and Pa_2 considerably reduce the number of performed rebalancing operations in most scenarios, whereas Pa_3 is the most effective strategy to improve the service level at the stations. Hence, the choice of the prioritization strategy to implement should depend on the objective pursued by the operator. The fourth strategy, Pa_4, also proved to be an effective alternative to improve service level while rebalancing proximal stations, with the potential to decrease operational routing costs.

Our strategies also seem particularly useful during periods in which the demand changes its patterns. The proposed prioritization strategies seem to adapt better to the demand changes in 2020 in comparison to BIXI's prioritization strategy. This outcome is due to the fact that the first is based on demand prediction while the former is based mostly on the location of the stations.

In future research, our strategies may be used along with routing optimization models, such that the entire rebalancing process is optimized in an integrated manner, taking into consideration the capacity of the dispatched trucks, as well as the loading/unloading of bikes among the stations.

CHAPTER 4 ARTICLE 2: TOWARDS EFFECTIVE REBALANCING OF BIKE-SHARING SYSTEMS WITH REGULAR AND ELECTRIC BIKES

Authors: Maria Clara Martins Silva, Daniel Aloise e Sanjay Dominik Jena.

Maria Clara Martins Silva: Development and implementation of the proposed model and simulation, experiments, validation, trip data analysis, literature review, writing and results analysis. **Daniel Aloise:** Development of the proposed model, writing, results analysis and supervision. **Sanjay Dominik Jena:** Development of the proposed model, writing, results analysis and supervision.

Under first round of revision at *Transportation Engineering* [1]

Date: August 1st, 2023

Abstract. The emerging demand for electric bicycles in recent years has prompted several Bike-Sharing Systems around the world to adapt their service to a new wave of commuters. Many of these systems have incorporated electric bikes into their network while still maintaining the use of regular mechanical bicycles. However, the presence of two types of bikes in a Bike-Sharing network may impact how rebalancing operations should be conducted in the system. Regular and electric bikes may exhibit distinct demand patterns throughout the day, which can hinder efficient planning of such operations. In this paper, we propose a new model that provides rebalancing recommendations based on the demand prediction for each type of bike. Additionally, we simulate the performance of our model under different scenarios, considering commuters' varying inclination to substitute their preferred bike with one of a different type. Our empirical experiments indicate the potential of our model to improve user satisfaction, reducing the total lost demand by approximately 10%, while reducing the lost demand for electric bikes by around 30%, on average, when compared to the existing rebalancing strategy used by the real-world Bike-Sharing System under study. Remarkably, this was accomplished while maintaining an almost identical average hourly count of rebalancing operations.

Keywords. *Bike-sharing, rebalancing, e-bikes, inventory management*

[1] available at [81]

4.1 Introduction

In the last years, we have seen bike-sharing systems (BSS) gaining the spotlight in the transportation scene due to their numerous advantages such as the absence of greenhouse gas emissions, the promotion of a healthy lifestyle, as well as easy and facilitated access. The history of this mobility service dates back to the 60s and has continuously evolved over time [111]. The fourth BSS generation, which we are currently experiencing, is marked by the inclusion of solar-powered docking stations, real-time system data, mobile apps, flexible parking, and electric bikes, also known as e-bikes [59].

In comparison with regular bikes, e-bikes are faster, easier to ride, especially on hilly paths, and overall they cause less fatigue [55]. To make BSSs more attractive to a group of commuters who are mainly interested in these advantages, BSSs around the world have introduced e-bikes into their networks – e.g. BIXI (Montréal), Citi Bike (New York). Nonetheless, regular bikes still please loyal commuters who search for health benefits or for a cheaper transportation mode. In [134], it is shown that the introduction of e-bikes in BSSs, alongside regular bikes, contributed significantly to the augmentation of BSSs revenues. However, a network with two types of bikes indeed introduces new challenges at every step of the service's logistics.

This challenge is especially hard for dock-based BSSs where the docks at the stations must be shared by both types of bikes. As such, too many bikes of a given type may lead to lost demand of the second type, and vice-versa, given that the number of docks is limited at each station. Hence, at the moment of rebalancing the inventory of a station, it is important to dynamically determine the number of ideal available bikes of each type in the stations of the system to guarantee its effective service.

Figure 4.1 presents the hourly average number of bikes rented on BIXI-Montreal in July 2022, where we can observe that the demand for regular and electric bike trips bounces over the day. Figure 4.2 highlights the considerable variation in e-bike demand per station at BIXI. While we can identify areas with high demand, it is notable that high-demand stations can be found adjacent to low-demand stations. This shows that understanding bike demand at station level is a complex task – even more in the presence of a heterogeneous bike offer. Additionally, the demand for bikes has undergone a significant transformation in recent years, driven by the shifts in working habits brought about by the COVID-19 pandemic [47, 108].

Besides, reallocating bikes through rebalancing operations can be quite expensive since it involves fuel costs, truck maintenance, driver's salary, etc. The trucks used for rebalancing are also responsible for CO_2 emissions and other polluting gases, which contradicts the BSS's

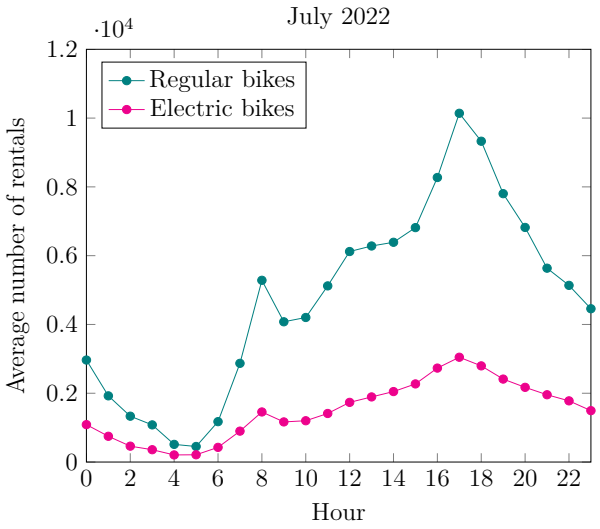

Figure 4.1 Hourly average number of rentals on BIXI-Montreal BSS in July 2022

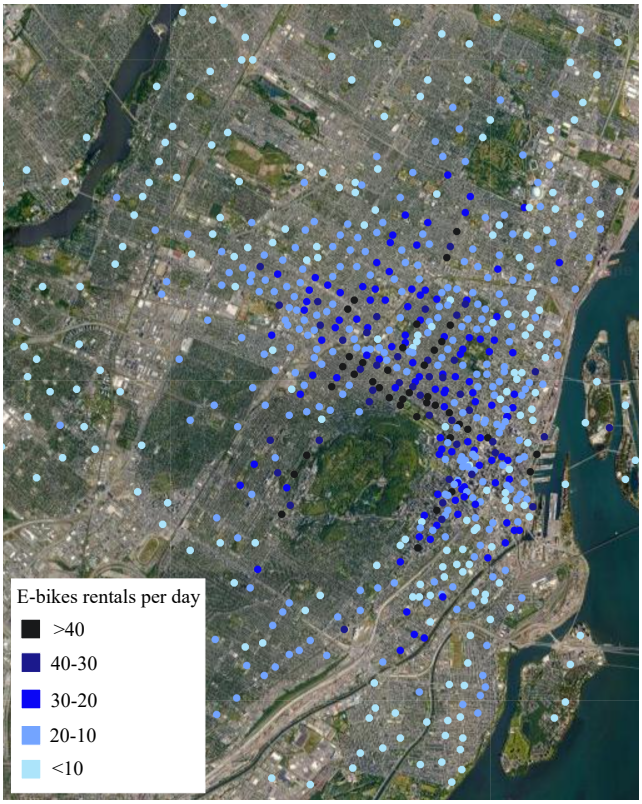

Figure 4.2 Average number of rented e-bikes per day at BIXI in July 2022

commitment to sustainability. Therefore, it is important to optimize the effectiveness of rebalancing operations, updating the inventories of stations that yield minimal lost demand

in the system. Nonetheless, leaving to the operator alone the task of understanding both demands and their correlations, while assuring that the rebalancing operations respect the BSS's limited resources, may lead to suboptimal rebalancing decisions.

The primary objective of our paper is to underscore the significance of incorporating demand predictions for both regular and electric bikes when making dynamic rebalancing decisions within a station-based bike-sharing system (BSS). In light of this, we propose a model that identifies imbalanced stations and determines the target inventory for each bike type. Essentially, our model provides recommendations for when and how many bikes of each type should be added or removed during the rebalancing process. By considering the expected demand for each bike type within a specific time period and accounting for station capacity, our model aims to optimize the rebalancing performance.

To evaluate our model effectiveness, we collected data from BIXI-Montreal[2] and conducted simulations to compare the inventory response between our proposed rebalancing strategy and the one currently employed by BIXI. Furthermore, our simulations explore various policies regarding the option of replacing one type of bike with another for a trip.

The paper is organized as follows. Section 4.2 reviews relevant literature on our research topic. Section 4.3 describes the proposed model to compute inventory intervals and target inventory for both regular and electric bikes. Section 4.4 presents the data, the tuning process, the inventory simulator, and the results of our experiments. Finally, our final remarks are given in Section 4.5.

4.2 Related works

One key performance indicator to assess a BSS is its *service level*, which is computed as the ratio between the number of satisfied trips and the total number of demanded trips [106]. It indicates whether a BSS is able to meet the commuters' demand, which is paramount for customer satisfaction.

The strategies to improve the service level in BSSs can be grouped into two main categories: network planning and operational rebalancing. The first consists of designing an ideal network configuration by optimizing the number of stations, docks and bikes, the location of the stations, initial inventories, etc., to reduce the total lost demand in the system [21, 97, 106]. The second represents an intervention, that can be performed either by the BSS operator or by the commuters, to redistribute the network's assets, e.g. bikes or batteries, among the stations [10, 32, 73, 93, 114].

[2] www.bixi.com

Table 4.1 Summary of strategies to improve the service level in EBSS.

Strategy	Research article	Network	Parking	Methodology
Operational rebalancing	Fukushige et al. [35]	Unimodal	Free-floating	User-based approach for rebalancing
	Tan et al. [114]	Unimodal	Parking locations	Optimization of routes for battery exchange
	Zhou et al. [133]	Unimodal	Free-floating	Optimization of routes for battery exchange
Network planning	Zhu [134]	Bimodal	Station-based	Optimization of bikes and e-bikes fleet
	Chen et al. [18]	Unimodal	Station-based	Optimization of bikes fleet and number of docks
	Zhou et al. [132]	Unimodal	Parking locations	Optimization of parking locations
	Martinez et al. [79]	Bimodal	Station-based	Optimization of stations locations
	Hosseini et al. [48]	Unimodal	Station-based	Optimization of stations performance
	Soriguera et al. [113]	Unimodal	Station-based	Optimization of e-bikes fleet, number of stations, number of docks, and rebalancing rate

Many studies in the literature have proposed strategies to improve the service level in BSSs. However, works that consider systems with shared e-bikes, hereafter denoted EBSSs, have just recently emerged. Table 4.1 summarizes the main works in the literature whose goal is to improve the service level in EBSSs. They are classified according to the strategy to improve the service level (operational rebalancing or network planning), the BSS network composition (modal or bimodal), the parking configuration (capacitated station-based, uncapacitated parking locations or free-floating), and their methodology.

Fukushige et al. [35], Tan et al. [114] and Zhou et al. [133] propose operational rebalancing strategies since they both deal with the reallocation of resources in the system. Fukushige et al. [35] study a user-based approach to better understand in which scenarios BSS commuters are stimulated to return bikes to a desired location under financial incentives. Tan et al. [114] and Zhou et al. [133] present models that propose travelling routes for exchanging discharged batteries for charged batteries among e-bikes in the system. Battery recharging is directly related to the performance of the service provided since e-bikes with discharged batteries remain unused in the system. In our work, we assume that battery recharging is conducted independently of the rebalancing process. This assumption is based on information provided by BIXI operators. At Bixi-Montreal, almost 20% of the stations are powered, allowing battery recharging on site.

The remaining references in the table address network planning strategies, i.e., they approach the problem of fleet dimensioning [18, 113, 134], dock dimensioning [18, 113]), locating the stations [79, 132], rebalancing rate, i.e., reallocated bikes per hour [113], and the correlation between the performance of the stations and external factors, such as the weather, population

characteristics and availability of nearby public transportation [48].

Our proposed model proposes to optimize the service level of a bimodal and station-based EBSS using target inventory values and inventory intervals to assist in the rebalancing process. The target inventory value represents the ideal number of bikes for a station and it is often used to establish how many bikes that station should have in a given time period in order to improve its performance [21, 45, 50, 53, 54, 57, 97]. Works proposing the computation of target inventory values vary depending on the metric used to assess the stations performance. For example, the performance of a station can be evaluated based on the expected lost demand [57, 97], the expected satisfied demand [45, 50, 53] or the commuters waiting times [21]. The inventory interval consists of an acceptable range in which the inventory can fluctuate while still meeting the expected demand. They are usually used to select which stations from the network need to be rebalanced [7, 53, 106].

4.3 Proposed model

In this Section, we present the model developed to automatically generate target inventory values and inventory intervals for bimodal BSS based on demand prediction.

4.3.1 Target inventory values

Target values refer to the ideal fill level to which the operator may want to set the station inventory for each bike type during the rebalancing process. Addressing the challenge of generating target values for two different types of demand that share station docks requires careful consideration to ensure that both demands are met without exceeding the station's capacity. So, to ensure the feasibility of the rebalancing recommendations provided by our target inventory values, our model divides the number of docks for each demand based on their respective service levels. Additionally, it establishes the optimal initial inventory within the allocated number of reserved docks for each demand.

Our model starts by training a machine-learning model to predict hourly rentals and returns at station-level for both types of bikes. Several studies in the literature focus on demand prediction of BSSs [see, e.g. 6, 15, 38, 66, 74, 76, 83]. In this work, we use a predictive model based on a Gradient-boosted tree, introduced in [53] which considers historical data as well as exogenous features such as weather conditions or holidays.

After forecasting the demand at each station of the BSS, we calculate their respective service levels. In our study, considering the availability of two types of bikes, we chose to compute the proportion of satisfied trips independently for each bike type. This approach allows us

to assess the service level for each type of bike individually, taking into account their specific demand patterns and availability.

To calculate the service level of a station, we model its inventory as a queue with a single server. The capacity of this queue is set to match the number of available docks at that particular station. Similar to other works that address the random rentals and returns of commuters in a BSS, such as [40, 43, 53, 60, 98, 106, 109], we assumed that the trips follow a Poisson distribution, so that the times between rentals and returns follow exponential distributions.

For a station s with an initial inventory of f and a specific number of docks, denoted as $C^{\mathcal{R}}$, allocated for regular bikes out of the total capacity of C_s docks, the expected service level for regular bikes during the time period $[0, T]$ can be computed as follows:

$$SL_s^{\mathcal{R}}(f, T, C^{\mathcal{R}}) = \frac{\int_0^T \mu_s^{\mathcal{R}}(t)(1 - p_s^{\mathcal{R}}(f, 0, t)) + \lambda_s^{\mathcal{R}}(t)(1 - p_s^{\mathcal{R}}(f, C^{\mathcal{R}}, t))dt}{\int_0^T \mu_s^{\mathcal{R}}(t) + \lambda_s^{\mathcal{R}}(t)dt}, \tag{4.1}$$

where $p_s^{\mathcal{R}}(f, N, t)$ is the probability that the station s stores N regular bikes at hour t, knowing that its initial inventory is equal to f at time 0; $\mu_s^{\mathcal{R}}(t)$ and $\lambda_s^{\mathcal{R}}(t)$ represent the predicted rental and return for regular bikes at hour t and station s. Here, the superscript $\mathcal{R}$ refers to values that are specific to regular bikes. Likewise, the service levels for e-bikes are computed as in eq. (4.1) by replacing $\mu_s^{\mathcal{R}}(t)$, $\lambda_s^{\mathcal{R}}(t)$, $p_s^{\mathcal{R}}(f, N, t)$ and $C^{\mathcal{R}}$ by $\mu_s^{\mathcal{E}}(t)$, $\lambda_s^{\mathcal{E}}(t)$, $p_s^{\mathcal{E}}(f, N, t)$ and $C^{\mathcal{E}}$, respectively, where the superscript $\mathcal{E}$ refers to values regarding e-bikes only.

Indeed, equation (4.1) depends on the number of docks $C^{\mathcal{R}}$ allocated to regular bikes at the analyzed time period. This allocation may vary to optimize the performance of the system based on the anticipated trip demand. In our model, the number of docks allotted for regular and electric bikes at station s for time period $[0, T]$, denoted $C_s^{\mathcal{R}}(T)$ and $C_s^{\mathcal{E}}(T)$, respectively, are determined as:

$$C_s^{\mathcal{R}}(T) = \underset{\overline{C} \in \{0, \ldots, C_s\}}{\arg\max} \left\{ \Lambda_s^{\mathcal{R}}(T, \overline{C}) + \Lambda_s^{\mathcal{E}}(T, \overline{C} - C_s) \right\}, \tag{4.2}$$

and

$$C_s^{\mathcal{E}}(T) = C_s - C_s^{\mathcal{R}}(T), \tag{4.3}$$

where function $\Lambda_s^{\mathcal{R}}(T, x)$ (resp. $\Lambda_s^{\mathcal{E}}(T, x)$) is chosen as the maximum or the average value of $SL_s^{\mathcal{R}}(f, T, x)$ (resp. $SL_s^{\mathcal{E}}(f, T, x)$) for $f \in \{0, \ldots, x\}$. The choice of the function as max or avg. influences the service level we want to optimize (the best or the average-case, respectively).

Once the number of docks reserved for regular bikes and e-bikes are determined, our model

proceeds to compute the target inventory values for regular and e-bikes for time period $[0, T]$ as:

$$\mathcal{T}_s^{\mathcal{R}}(T) = \underset{f \in \{0,\dots,C_s^{\mathcal{R}}(T)\}}{\arg\max} \{SL_s^{\mathcal{R}}(f, T, C_s^{\mathcal{R}}(T))\}, \tag{4.4}$$

and

$$\mathcal{T}_s^{\mathcal{E}}(T) = \underset{f \in \{0,\dots,C_s^{\mathcal{E}}(T)\}}{\arg\max} \{SL_s^{\mathcal{E}}(f, T, C_s^{\mathcal{R}}(T))\}. \tag{4.5}$$

The above equations compute the target values as the inventory fill levels that maximize the service level for the corresponding bike type. Note again that the number of docks allotted to regular and electric bikes, $C_s^{\mathcal{E}}(T)$ and $C_s^{\mathcal{E}}(T)$, respectively, sum to the total dock capacity of the station, as defined in Equation (4.3). As such, the computed target values respect the capacity of the station and the number of docks reserved for each type of bike. This ensures that the recommended rebalancing actions based on the target inventory values are practical and feasible to implement.

4.3.2 Inventory intervals

Inventory intervals typically serve as an indicator when a station should be rebalanced: this is the case when the station inventory falls outside of the defined interval. Considering only the total rentals or returns without distinguishing between different types of demand can obscure the identification of lost demand for a specific type of bike. Similarly, creating inventory intervals tailored to each demand while assuming that the total number of docks is always available can lead to undesirable situations. For instance, a station may become completely full without triggering any alerts because neither demand (for regular or e-bikes) has exceeded its upper or lower bounds. Therefore, it is crucial to take into account the station capacity when setting inventory intervals to ensure optimal inventory management and prevent potential issues.

To calculate the inventory intervals, we begin by computing the maximum and minimum service levels for each demand for time period $[0, T]$, which are given by

$$\underline{SL}_s^{\mathcal{R}}(T) = \min_{f \in \{0,\dots,C^{\mathcal{R}}\}} SL_s^{\mathcal{R}}(f, T, C_s^{\mathcal{R}}(T)), \tag{4.6}$$

and

$$\overline{SL}_s^{\mathcal{R}}(T) = \max_{f \in \{0,\dots,C^{\mathcal{R}}\}} SL_s^{\mathcal{R}}(f, T, C_s^{\mathcal{R}}(T)) \tag{4.7}$$

for regular bikes. These values can be analogously obtained for e-bikes, by replacing the

superscript $\mathcal{R}$ by $\mathcal{E}$.

Then, accepted service levels at station s for time period $[0, T]$ are calculated for each type of bike as:

$$\Omega_s^{\mathcal{R}}(T) = \underline{SL}_s^{\mathcal{R}}(T) + \beta^{\mathcal{R}}(\overline{SL}_s^{\mathcal{R}}(T) - \underline{SL}_s^{\mathcal{R}}(T)), \tag{4.8}$$

and

$$\Omega_s^{\mathcal{E}}(T) = \underline{SL}_s^{\mathcal{E}}(T) + \beta^{\mathcal{E}}(\overline{SL}_s^{\mathcal{E}}(T) - \underline{SL}_s^{\mathcal{E}}(T)), \tag{4.9}$$

The model incorporates two hyperparameters, $\beta^{\mathcal{R}}$ and $\beta^{\mathcal{E}}$, which are specific to each type of bike. These hyperparameters provide flexibility for the operator to fine-tune the computed inventory intervals for a station based on the behaviour patterns of its user base. By separately adjusting the values of $\beta^{\mathcal{R}}$ and $\beta^{\mathcal{E}}$, the operator can also customize the inventory intervals to be more or less stringent for each demand type throughout the day.

Finally, the inventory intervals for regular bikes and e-bikes at station s for time period $[0, T]$ are computed as:

$$\mathcal{I}_s^{\mathcal{R}}(T) = \{f \in \{0, ..., C_s^{\mathcal{R}}(T)\} | SL_s^{\mathcal{R}}(f, T, C_s^{\mathcal{R}}(T)) \geq \Omega_s^{\mathcal{R}}(T)\} \tag{4.10}$$

and,

$$\mathcal{I}_s^{\mathcal{E}}(T) = \{f \in \{0, ..., C_s^{\mathcal{E}}(T)\} | SL_s^{\mathcal{E}}(f, T, C_s^{\mathcal{E}}(T)) \geq \Omega_s^{\mathcal{E}}(T)\} \tag{4.11}$$

Remark that, by definition, inventory intervals contain target inventory values. This avoids that a station could be categorized as unbalanced immediately after a rebalancing operation has taken place.

Finally, because the availability of e-bikes is compromised due to battery discharging, there might be an underestimation of the actual demand for e-bikes used in training our Gradient-boosted tree. In response to this concern, we increased the predicted hourly demand for e-bikes by 10%, aligning with observations made by BIXI operators who reported us that, on average, approximately 10% of e-bikes are uncharged per hour in the system. Thus, we aim to obtain a more accurate computation inventory intervals and target inventory values, by identifying additional rebalancing demand for e-bikes.

4.4 Computational Experiments

In this section, we assess our model in comparison with the approach used by BIXI in 2022. First, we will present the data used in our experiments. Then, we briefly explain our

simulations to emulate the inventories based on the rebalancing strategy applied. Next, we discuss the process for selecting the best hyperparameters, $\beta^{\mathcal{R}}$ and $\beta^{\mathcal{E}}$, for our model. At last, we present the results collected from our experiments.

4.4.1 Data

In light of the fact that BIXI added a considerable amount of e-bikes to its network in 2022, we opted to collect only the data from the aforementioned year. Thus, the data used in our experiments contain hourly information from April to September 2022, being grouped into three categories: temporal, weather, and trip data. The first category includes time features, such as hour, day, day of the week, month, and holidays. The second category contains data describing the weather, such as temperature, humidity, rain, and wind speed. Both temporal and weather data were collected from the official website of the Government of Canada[3] (except for the holiday feature which was manually noted). The trip data is composed of the number of rentals and returns at each station and it was provided by BIXI[4]. In addition to the data mentioned before, BIXI also provided the inventory intervals used in 2022 and network information that includes the capacity of the 745 stations and the proportion of regular bikes ($\approx 75\%$) and electric bikes ($\approx 25\%$) in their network.

The collected data was divided between train, validation, and test datasets. Given that we have access to the inventory intervals manually computed by BIXI operators for August and September 2022, we assign these months to the validation and test dataset, allocating the first 15 days of August and September to the validation dataset and the remaining days of these months to the test dataset. This distribution was chosen to evaluate the rebalancing strategies under varying weather conditions, as these months can present contrasting temperatures and rain features. The training dataset, meanwhile, contains data from April to July 2022.

The proportion of trips made with regular bikes and e-bikes is consistent across all datasets. In the training dataset, approximately 76% of trips were made using regular bikes whereas 24% of the trips used e-bikes. In the validation and test datasets, the proportion of trips using regular bikes and e-bikes is 75% and 25%, respectively.

4.4.2 Experiment

Our proposed model, denoted hereafter shared-$\mathcal{RE}$, is compared against a baseline approach, namely B0, that corresponds to the strategy applied by BIXI operators in 2022. The inventory

[3]https://climate.weather.gc.ca/
[4]https://bixi.com

intervals and target inventory values used to assist the rebalancing operations at BIXI were manually determined by their operators. These decisions were based on historical trip data at each station, without differentiation between regular bikes and e-bikes.

Simulation of B0

Simulation B0 begins by initializing the inventories of regular and electric bikes at the stations with their respective target inventory values. However, it is worth noting that BIXI employs a unique target inventory value for each station during specific time periods. To decouple this single target inventory value between regular bikes and e-bikes, a straightforward approach is to distribute it based on the proportion of each bike type rented at the station during the observed period. This distribution can be determined by analyzing the training dataset, which ensures that the distribution of bikes aligns with the observed rental patterns.

Then, at each simulated hour, the inventory of each station is updated with the historical rentals and returns from BIXI data. At this point, the simulation verifies which stations trigger a rebalancing alert, that is, which stations surpass the bounds of their inventory interval. The simulation also keeps track of the lost demand, i.e., how many bikes were missing during the rentals and how many bikes could not be returned due to full stations by assuming that all rentals and returns happen simultaneously at every simulated hour. Thus, we stress the network to capture its possible failures. We note that uncharged bikes still occupy docks for our simulation.

Due to constrained rebalancing capacity, the alerted stations are prioritized based on their degree of imbalance, computed as the deviation of their inventory from the target value. Consequently, only a limited subset of stations, within the capacity threshold, undergo rebalancing procedures per hour.

Figure 4.3 illustrates the simulation of the bike inventory of a station using B0. In the illustration, the station raised an alert due to the shortage of bikes, i.e. its inventory (3) is below the inventory lower bound (4). Assuming that the station is selected to be rebalanced, the inventory of regular and e-bikes is updated according to the target inventory value and the station's historical demand observed in the training data.

Simulation of our model

shared-$\mathcal{RE}$ provides inventory intervals and target inventory values for each type of bike used in the EBSS. Therefore, alerts are individually raised for each type of bike demand, and stations are rebalanced according to the associated target inventory values.

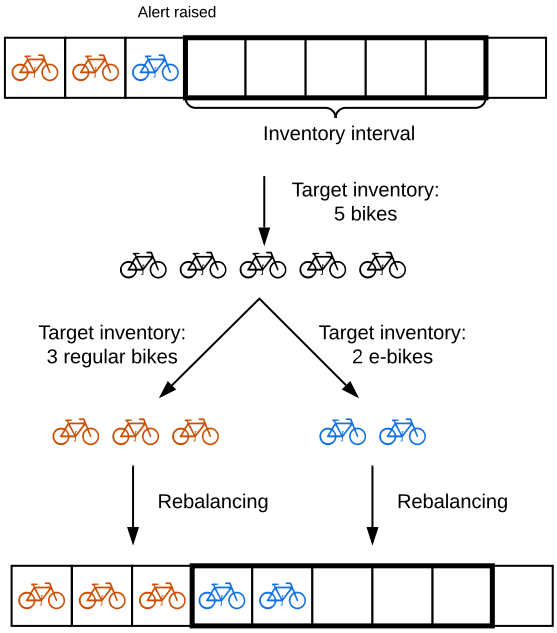

Figure 4.3 Simulation of the inventory and rebalancing process in the model B0. Orange bikes represent regular bikes whereas blue bikes represent e-bikes.

Similar to the simulation of B0, the simulation of shared-$\mathcal{RE}$ emulates the inventory based on rentals, returns, and rebalancing operations conducted hourly accounting for the rebalanced capacity of the system. However, in this scenario, the target inventory and inventory intervals are customized for each demand. This allows the rebalancing alerts to identify and address inventory deficiencies or excesses specific to each demand, and calculate target inventory values accordingly.

Since the inventory intervals generated by shared-$\mathcal{RE}$ are associated with each type of demand, it is expected that a greater number of alerts will be raised, potentially resulting in increased rebalancing activities. To address this concern, our model restricts the hourly rebalancing operations based on the recorded rebalancing operations of B0. Consequently, if the number of alerts raised by shared-$\mathcal{RE}$ within an hour exceeds the number of rebalancing operations performed in that same hour in the B0 simulation, we select stations at random for rebalancing among those that raised alerts.

During the simulation of our model, the rebalancing process is conducted to replenish the inventory of a station, taking into account the target values for regular and electric bikes. These target values are separately computed based on the demand for each type of bike, as

explained in Section 4.3.2. This approach enables the rebalancing process to be triggered by the demand of a specific bike type. In Figure 4.4, for example, only the inventory of e-bikes drops below the lower bound of its inventory interval. Subsequently, the rebalancing process focuses on restoring both the regular and e-bike inventories to their respective target values. This reflects a more realistic operational scenario, as an employee is already dispatched to the station for replenishment. In the given example, two e-bikes and one regular bike are added to the station during the rebalancing process.

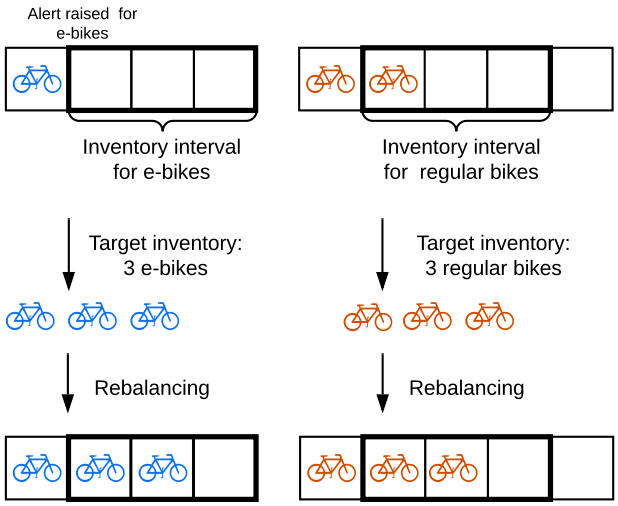

Figure 4.4 Simulation of the inventory and rebalancing process in our model. Orange bikes represent regular bikes whereas blue bikes represent e-bikes.

Additional simulation remarks

This section presents remarks that are applicable to both simulations detailed in previous sections.

First of all, the inventory of the stations is initialized at the beginning of the simulations by distributing the total number of bikes in the system (i.e., 6377 regular bikes and 2015 electric bikes for the considered simulated period) proportionally to the target inventory values of each bike demand at the first simulated hour. More precisely, the initial inventory of regular bikes at each station s is made equal to:

$$b_{regular} \times \frac{\mathcal{T}_s^{\mathcal{R}}(1)}{\sum_{s' \in S} \mathcal{T}_{s'}^{\mathcal{R}}(1)},$$

where $b_{regular}$ corresponds to the total number of regular bikes in the system and S to the set of stations. The initial inventory of e-bikes at each station is defined likewise.

Throughout the simulation, the amount of bikes in the system is kept consistent by employing an algorithm that selects stations while ensuring an equilibrium between the bikes taken from and added to these stations. A detailed explanation of this procedure is provided in A.

Finally, the simulations take into account that 10% of the electric bikes are uncharged per hour. These bikes are randomly chosen at the beginning of each simulated hour, making them unavailable in the system, and hence, incurring lost demand. The code for both simulations, as well as the code of the proposed model, can be accessed in the repository[5].

4.4.3 Analysis of commuters preferences

Based on the commuters' preferences, four different scenarios are emulated regarding bike substitutions:

- None: in this scenario, users never replace their desired bike with another type.
- All bikes: in this scenario, users are flexible in their preferences and will always accept any available bike, regardless of their initial choice.
- Reg. bike → E-bikes: only users who seek regular bikes accept an e-bike if the first is unavailable.
- E-bike → Reg. bikes: only users who seek for an e-bike accept a regular bike if the first is unavailable.

By simulating these different scenarios, we can analyze the impact of bike substitution preferences on the overall bike availability and system performance. This provides insights into the feasibility and desirability of allowing bike substitutions in a bike-sharing system and helps optimize the allocation and utilization of bikes based on customer preferences.

4.4.4 Tuning

The values of $\beta^{\mathcal{R}}$ and $\beta^{\mathcal{E}}$ were tuned through simulations with the validation set with the objective of minimizing the lost demand – recall that the number of rebalancing operations

[5] `https://github.com/datascientistbss/Paper_Journal`

are limited by B0. Table 4.2 presents the optimized hyperparameter values, where shared-$\mathcal{RE}_{max}$ and shared-$\mathcal{RE}_{avg}$ refer to the use of $\Lambda_s(\cdot)$ as the maximum or the average service level obtained for different values of the initial inventory (see section 4.3.1 - eq.4.2).

The analysis of Table 4.2 reveals a discernible pattern in the values of the hyperparameters $\beta^{\mathcal{R}}$ and $\beta^{\mathcal{E}}$ based on the bike substitution policy implemented. When there are no bike substitutions or restrictions on bike types, the values of both hyperparameters are fairly similar. However, in scenarios where regular bikes can be substituted with e-bikes, the inventory intervals display greater stringency towards e-bikes, leading to lower values of $\beta^{\mathcal{R}}$ and higher values of $\beta^{\mathcal{E}}$. Conversely, in the scenario where only electric bikes can be replaced by regular bikes, the hyperparameter results exhibit the opposite trend. This demonstrates the model's ability to prioritize each demand independently, as wells as its capacity to adapt to the users preferences.

Table 4.2 Optimized hyperparameters values used in the tests.

Bike substitution	Model	($\beta^{\mathcal{R}}$,$\beta^{\mathcal{E}}$)
None	shared-$\mathcal{RE}_{max}$	(0.4,0.3)
	shared-$\mathcal{RE}_{avg}$	(0.4,0.5)
All bikes	shared-$\mathcal{RE}_{max}$	(0.4,0.2)
	shared-$\mathcal{RE}_{avg}$	(0.4,0.4)
Reg. bikes → E-bikes	shared-$\mathcal{RE}_{max}$	(0.3,0.5)
	shared-$\mathcal{RE}_{avg}$	(0.3,0.5)
E-bikes → Reg. bikes	shared-$\mathcal{RE}_{max}$	(0.4,0.3)
	shared-$\mathcal{RE}_{avg}$	(0.4,0.3)

4.4.5 Results

Our results compile the number of rebalancing operations and lost demand computed from the simulations of the baseline and our proposed models. The inventory intervals and the target values for B0 were provided by BIXI for the following time periods in a day: from 6 am to 9 am, from 9 am to 11 am, from 11 am to 4 pm, from 4 pm to 7 pm, and from 10 pm to 6 am. These same time periods were used by shared-$\mathcal{RE}$ to compute inventory intervals and target values. Besides, the root mean squared error (RMSE) of the Gradient-boosted tree demand prediction on the test data was 1.70.

Our experiments considered an hourly rebalancing capacity of 50 stations. This value was provided to us by BIXI as the valid capacity for the year 2022.

Table 4.3 presents the results regarding the number of rebalancing operations per hour and the percentage of lost demand over the total served demand from the simulations of models B0, and shared-$\mathcal{RE}$ and its different settings of bike substitution.

Table 4.3 Results regarding the simulated models: the average number of rebalancing operations per hour and the lost demand (in % with respect to the total served demand).

Simulated model	**Bike substitution**	**Total lost demand %**	**Regular bikes lost demand %**	**E-bikes lost demand %**	**Rebalancing per hour**
B0	None	3.32	2.46	5.89	25.08
shared-$\mathcal{RE}_{max}$		3.07	2.75	4.01	25.77
shared-$\mathcal{RE}_{avg}$		3.00	2.67	3.98	26.40
B0	All bikes	2.25	1.84	3.45	23.16
shared-$\mathcal{RE}_{max}$		2.10	1.88	2.80	25.42
shared-$\mathcal{RE}_{avg}$		2.07	1.83	2.78	26.61
B0	Reg. bikes → E-bikes	2.89	1.86	5.99	25.63
shared-$\mathcal{RE}_{max}$		2.66	1.97	4.85	26.48
shared-$\mathcal{RE}_{avg}$		2.64	1.93	4.78	26.86
B0	E-bikes → Reg. bikes	2.98	2.84	3.42	22.36
shared-$\mathcal{RE}_{max}$		2.81	2.95	2.40	24.50
shared-$\mathcal{RE}_{avg}$		2.80	2.93	2.43	24.36

In summary, the results show that:

- The performance of both shared-$\mathcal{RE}_{max}$ and shared-$\mathcal{RE}_{avg}$ indicates a minimal difference in terms of lost demand and rebalancing operations when implementing the dock division at a station based on either the best or average service level provided.
- Our model consistently outperforms the baseline model, reducing the total lost demand by up to 10% (bike substitution = 'None', shared-$\mathcal{RE}_{avg}$). More specifically, the inventory intervals and target inventory values generated by our model have proven to be highly effective in decreasing lost demand for e-bikes when compared to BIXI's rebalancing strategy, reducing the lost demand for e-bikes by up to 32%. This demonstrates that shared-$\mathcal{RE}_{max}$ and shared-$\mathcal{RE}_{avg}$ effectively identify and adapt to the increasing demand for e-bikes better than B0.

- When comparing the various configurations of user preferences, the results align with our expectations across all models. The scenario where no bike substitution is allowed generates the highest lost demand, while the scenario where any substitution is accepted yields the lowest demand loss. The remaining scenarios fall somewhere in between these extremes. Notably, the scenario where only regular bikes can be replaced by electric bikes results in less lost demand than the reverse scenario. This can be attributed to the considerably higher demand for regular bikes observed in our simulation data, which can be partially attributed to the fact that 76% of BIXI's bikes are regular ones. Consequently, the majority of lost demand cases computed in our simulations involve regular bikes.

- By introducing the option to replace regular bikes with electric bikes, the occurrences of lost demand can be significantly reduced. This is due to the potential to fulfill the demand for regular bikes with available electric bikes, thereby mitigating lost opportunities for riders.

- Overall, B0 requires fewer rebalancing interventions than our model. This result is expected as the simulation shared-$\mathcal{RE}$ can detect imbalances of each type of demand, leading to a higher number of alerts and, consequently, a higher number of rebalancing activities. Nevertheless, the rebalancing operations carried by shared-$\mathcal{RE}$ respect BIXI's maximum rebalancing capacity.

- We observe that the average number of rebalancing operations per hour in all simulations is significantly lower than the hourly rebalancing capacity of 50 stations. This is attributed to the reduced demand typically observed from 10 pm to 7 am, as illustrated in Figure 4.1. During these off-peak hours, the number of alerts and, consequently, rebalancing operations, is considerably lower than 50. In contrast, in periods of high demand the number of alerts is often greater than the rebalancing capacity – in a single simulated hour, the maximum number of observed alerts was 398.

4.5 Conclusion

4.5.1 General discussion

Rebalancing bike-sharing systems is a multifaceted task that encompasses numerous factors such as demand variability, time sensitivity, and user preferences. As electric bikes become more popular in existing systems, accounting for such additional demand adds complexity to the rebalancing planning. We propose a model capable of providing targeted rebalancing

recommendations for station-based Electric Bike Sharing Systems, comprising both regular and e-bikes. Our model leverages predicted demand for the upcoming hours to tailor recommendations specific to the demand of each bike type.

From a theoretical perspective, our work extends existing works on the definition and computation of service-levels [see, e.g. 106], inventory target values and inventory intervals [see, e.g. 53] for regular bikes to bimodal systems with both regular and electric bikes. Here, a simultaneous computation is required to ensure that the total station capacity is not exceeded when considering both types of resources. Our model offers an automated division of docks per station based on predicted demand while allowing for customization according to the operator's requirements for each demand. One significant advantage is that our model independently adjusts the inventory intervals for each bike type. This flexibility is crucial, as it accommodates the varying preferences of BSS users. One proposed model variant, namely shared-$\mathcal{RE}_{max}$, computes the inventory intervals and target inventory values aiming to maximize the ratio of satisfied trips while dividing the available docks between regular and electric bikes. Conversely, the alternate variant, denoted shared-$\mathcal{RE}_{avg}$, aims to maximize the average of this ratio across varying initial inventory values. Both variants are compared against B0, a rebalancing strategy emulating the current practice at BIXI Montreal.

From a practical perspective, the mechanisms of inventory intervals and target inventory values are often already an essential part of the rebalancing process. A major advantage of our approach is hence its relative simple deployment in existing decision-making processes and its minimal requirement of resources: a computer with moderate processing power and a database to feed the prediction model. Furthermore, the manual computation of inventory targets and intervals, although common in BSSs, may result in suboptimal performance. This issue is particularly pronounced in the context of bimodal EBSSs, where the nuances of each demand type may be overlooked. Such manual computation can hence be easily replaced by an automatized computation of such parameters, without the need of restructuring the existing information and decision-making process.

Our empirical experiments show that our proposed model is able to reduce the amount of total lost demand in all simulated scenarios, demonstrating its ability to adapt to diverse trip patterns and commuters' preferences. In the scenario with the greatest difference in performance, our model managed to reduce the total lost demand by up to 10% and the lost demand for electric bikes by up to 32% over the baseline model B0.

The results further demonstrate the importance of comprehending commuters' preferences and their willingness to substitute their initial bike choice when designing a rebalancing strategy. This understanding enables operators to make informed decisions regarding the

supply of each bike type, ensuring the provision of a high-quality service. Additionally, our results show that actively encouraging commuters to consider alternative bikes when their desired option is unavailable can have a significant impact on reducing the lost demand. This effect was particularly pronounced when the initially preferred bike type exhibits higher demand compared to the other. By promoting bike substitution, operators can effectively mitigate the occurrence of lost demand, leading to improved service reliability and user satisfaction.

Finally, from a financial perspective, the predictive models used within our method make use of open-source libraries, therefore circumventing the need for licensing costly optimization solvers, which would be required when using classical rebalancing optimization models. However, it is important to note that our algorithm proposes rebalancing recommendations on a tactical level, i.e., it identifies the ideal station inventories without explicitly proposing vehicle routes. Operators are therefore required to design rebalancing routes based on these recommendations, either by manual planning or using routing algorithms.

4.5.2 Limitations and future work

Due to the use of actual trip data in our experiments, we have not considered information about trips that were not undertaken due to the unavailability of bikes or docks, also referred to as unobserved demand. In future research, we plan to address this limitation by conducting experiments using synthetic data. This will enable us to explore a wider range of scenarios and accurately quantify demand losses, allowing for a more comprehensive evaluation of the rebalancing recommendations.

Finally, the rebalancing recommendations provided by our model can be seamlessly integrated into optimization routing models. This integration would allow for the optimization of the entire rebalancing process in a unified manner, maximizing the effectiveness and efficiency of the system as a whole. This direction holds promise for future research and offers potential for further improvements in the field of bimodal bike-sharing system management.

CHAPTER 5 ARTICLE 3: DYNAMIC REBALANCING FOR BIKE-SHARING SYSTEMS UNDER INVENTORY INTERVAL AND TARGET PREDICTIONS

Authors: Jiaqi Liang, Maria Clara Martins Silva, Daniel Aloise e Sanjay Dominik Jena.

Jiaqi Liang: Development and implementation of DROB-T, DROB-I, data trip generator and simulation, experiments, literature review, writing and results analysis. **Maria Clara Martins Silva:** Development and implementation of the data trip generator, validation of the simulation, experiments, generation of inventory intervals, target inventory, demand prediction, and noisy demand prediction, data trip analysis, literature review, writing and results analysis. **Daniel Aloise:** Development of the DROB-T, DROB-I and data trip generator, writing, results analysis and supervision. **Sanjay Dominik Jena:** Development of the DROB-T, DROB-I and data trip generator, writing, results analysis and supervision.

Submitted to *Transportation Research - Part C*.

Date: November 14th, 2023

Abstract. Bike-sharing systems have become a popular transportation alternative. Unfortunately, station networks are often unbalanced, with some stations being empty, while others being congested. Given the complexity of the underlying planning problems to rebalance station inventories via trucks, many mathematical optimization models have been proposed, mostly focusing on minimizing the unmet demand. This work explores the benefits of two alternative objectives, which minimize the deviation from an inventory interval and a target inventory, respectively. While the concepts of inventory intervals and targets better fit the planning practices of many system operators, they also naturally introduce a buffer into the station inventory, therefore better responding to stochastic demand fluctuations. We report on extensive computational experiments, evaluating the entire pipeline required for an automatized and data-driven rebalancing process: the use of synthetic and real-world data that relies on varying weather conditions, the prediction of demand and the computation of inventory intervals and targets, different reoptimization modes throughout the planning horizon, and an evaluation within a fine-grained simulator. Results allow for unanimous conclusions, indicating that the proposed approaches reduce unmet demand by up to 34% over classical models.

Keywords. *Bike-sharing systems, dynamic rebalancing, inventory Intervals, target inventories, reoptimization modes, mixed-integer programming*

5.1 Introduction

The pursuit of environmentally friendly transportation modes has increased considerably in the last few years, with BSSs having emerged as a notable choice. As of 2022, there were over 1,900 BSSs in operation comprising almost 9 million bikes [87], offering cities opportunities to reduce carbon emissions, traffic relief, and improve the quality of life for their residents [19].

A particular challenge in the management of BSSs are station imbalances. Most users follow consistent travel patterns, often commuting toward commercial areas (such as city centers) during morning peak hours and returning to residential zones after work. Such behavior often results in stations being full or empty, leading to *lost demand*, i.e., rental demand that cannot be met due to an empty station or return demand that cannot be met as a result of a full station. Occasional user trips introduce stochasticity, further aggravating station imbalances. As a remedy, BSSs operators often redistribute bikes among the stations, typically, via trucks, a process known as *rebalancing*. Rebalancing imbalanced stations has been proven to be more cost-effective than alternative solutions, such as adding more stations or installing additional docks [109]. The development of effective rebalancing strategies has therefore become a crucial research field with the potential to significantly improve user satisfaction. Two primary rebalancing schemes have been acknowledged in the literature: overnight station rebalancing and intraday rebalancing. The latter is often referred to as the Dynamic Bicycle Repositioning Problem (DBRP) [88, 98]. In contrast to overnight rebalancing, the DBRP involves continuous intraday rebalancing operations in parallel to the user trips occuring throughout the day. Given its higher impact on demand satisfaction, we here focus on this problem.

While a few recent works use Markov Decision Processes (MDP) to address the problem variants [9, 64, 65, 69, 107], the majority of the literature applies Mixed-integer Programming (MIP) models [44, 82, 129, 131], given that they are flexible and widely used within industrial decision-making processes. Among MIP models, multi-period models benefit from an integrated planning over all time-periods and do therefore not suffer from the myopic behavior of single-period models.

Most models aim at minimizing lost demand or maximizing successful trips [44, 73, 110, 128, 129, 130], with both rental and return demands estimated either by naive predictions (e.g., the historical mean) [41, 43, 73, 129] or more sophisticated Machine Learning (ML) techniques [101, 130]. However, minimizing lost demand may result in sub-optimal rebalancing solutions if the predicted rental and return are not accurate enough. BSSs operators, like

BIXI Montreal, often use *inventory intervals* and *target inventories* within their rebalancing planning. The inventory interval of a station refers to the acceptable range of its inventory, ensuring that the station maintains a level of available bikes and free docks. The target inventory of a station represents the desired number of bikes at that station to ensure optimal service.

While many methods have been proposed to compute inventory intervals and target inventories (see e.g., [50, 53, 72, 97]), only a few have incorporated them into optimization models. Notably, most relevant studies [45, 62, 106] either focus on single-period models or minimize the deviation of target inventory at the end of the planning process only. Although [118] and [120] attempt to introduce intervals into multi-period models, they relax the intervals by using station capacities directly in their experiments. Furthermore, the intervals and targets in these models are often determined without considering weather conditions, which significantly impact user behavior. As a result, the benefits of combining optimization algorithms with intervals or targets in the objective function for multi-period models are still to be determined.

In a similar vein, while some of the existing models have been carried out in a rolling fashion (see, e.g., [44, 45, 82, 101, 110]), literature has not yet quantified the benefits of integrating system status update through reoptimization (i.e., rolling and folding planning) over classical static planning of multi-period rebalancing models. In this paper, we aim at filling these gaps and approach the DBRP by incorporating inventory intervals and target inventories.

Contributions. Our work evaluates the entire pipeline required for an automatized and data-driven rebalancing process in BSSs. The main contributions can be summarized as follows. (i) We propose two optimization models that integrate inventory intervals and target inventories into the objective functions, concepts that are often already used within the decision-making process of BSS operators. In contrast to classical models that minimize unmet demand, the proposed models tend to ensure a buffer in the station inventories and are therefore more capable of dealing with demand fluctuations. (ii) We propose a realistic instance generator, generating varying weather conditions for different days along with trip data that is historically coherent with such conditions. (iii) We conduct an extensive comparison among three multi-period models with different objective functions for DBRP, including the classical objective that minimizes unmet demand and the two proposed objectives that minimize the deviations from inventory intervals and target inventories. Demand predictions are obtained from an advanced machine learning model, capable of making sufficiently accurate predictions based on weather and temporal features. Inventory and target inven-

tories are computed such that they maximize the desired service-level. The performance is estimated by a fine-grained discrete-event simulator. Our models demonstrate a remarkable robustness to cope with trip fluctuations, reducing lost demand by up to 34% as compared to the model minimizing lost demand. (iv) We empirically compare the impact of employing different reoptimization modes (i.e., static and rolling planning) for all models. The results indicate a clear advantage of reoptimizing over the planning horizon, reducing the lost demand by at least 30% on average, without necessarily increasing the computing time. (iv) We compare the impact of using perfect information and less accurate demand predictions on the performance of the planning models. Interestingly, our proposed optimization models remain remarkably robust. (v) A case-study on real-world data is considered, confirming the benefits of the proposed approaches.

Outline. This paper is organized as follows. Section 5.2 reviews relevant literature for BSSs in objective functions in rebalancing models, inventory intervals and target inventories, trip prediction, and reoptimization modes. Section 5.3 reviews the baseline model that minimizes unmet demand and introduces two dynamic rebalancing models minimizing the deviations from inventory intervals and target inventories. Numerical experiments and analysis on synthetic and real-world data are presented in Section 5.4. This is followed by the conclusions in Section 5.5.

5.2 Literature Review for Rebalancing Problems in BSSs

This section reviews the literature related to the here considered planning problem and our contributions, focusing on the objective functions used within DBRP models, inventory intervals and targets, demand prediction, and reoptimization modes.

5.2.1 Objective functions in rebalancing models

We first review existing MIP models and their objective functions used in dynamic rebalancing. Metrics used in the objective functions can be classified into three different types (see Appendix of [70]): distance-based metrics, loading-based metrics, and demand-based metrics. Distance-based metrics are associated with the travelling distance of vehicles, mainly including travelling costs, travelling time, and fuel consumption (see e.g., [1, 43, 95, 131]). Loading-based metrics are associated with the number of handling (i.e., loading and unloading) operations (see e.g., [49, 115]). Handling cost or time reflects the workload of operations. Finally, demand-based metrics concern the dissatisfaction of customers. Some studies con-

sider more than one aspect in their objective functions (see, e.g., [41, 49, 62, 82, 130]).

We here focus on demand-based metrics, which have been more relevant in the literature, and to which our contributions are directly related. The most common approach for demand-based metrics is to minimize the lost rental and return demand, which is also equivalent to maximizing successful trips (see e.g., [20, 43, 44, 49, 73, 110, 128, 129, 130, 131]). Concurrently, there have been a few attempts to minimize the deviation between the inventories of the stations and their target inventories [45, 62] or inventory intervals [120], which still holds considerable potential for further exploration. Our work focuses on such inventory intervals and target inventories, for which the related literature is reviewed next.

5.2.2 Inventory intervals and target inventories

As opposed to minimizing lost demand, the concepts of inventory intervals and target inventories have been found to be useful within the rebalancing decision-making process. Inventory intervals define an acceptable range of the bike inventory at individual stations, whereas target inventories represent specific inventory levels that operators aim to uphold at each station. Both are typically designed to ensure a high rate of demand satisfaction.

Even though inventory intervals and target inventories are often used concepts in the planning processes of BSS operators, only a few works have incorporated them into optimization models. When using target inventories, objective functions typically minimize the deviations between the station inventories and the specified targets. Here, [45] consider a single-period model, executed in a rolling planning, that minimizes such inventory deviations. Further, [62] propose a multi-period rebalancing model for which they compare various metaheuristics. Next to multiple criteria (such as minimizing costs) considered within the objective function, the model also aims at aligning the final station inventory at the end of the planning horizon with a specific target inventory. Finally, the model introduced by [13] also uses target values, but instead of minimizing deviations, they use hard constraints to ensure that station inventories equal those targets at the end of the rebalancing process. As a result, the minimization of the deviations between station inventories and target values at each of the time-periods, as a mean of improving demand satisfaction, has not yet been considered and deserves further investigation.

Inventory intervals have also been implemented into rebalancing models as hard constraints [106, 118], requiring that station inventories remain within a specific range. Here, [106] focus on the static rebalancing problem with a single time-period, while [118] propose a multi-period model. However, the authors do not consider these constraints in their computational experiments. Given that hard constraints may easily lead to infeasible outcomes, [120] present

a multi-period model that, among several other criteria, minimizes the deviation of station inventories from predefined inventory intervals within the objective function. Unfortunately, their experimental setup also ignores such intervals, leaving the actual effectiveness of inventory intervals within the optimization model unexplored.

While the works cited above assume that target inventories and inventory intervals are given, a few works also propose how to effectively compute them. Target inventories have often been computed such that they reduce the probability of a station reaching both the full and empty status, typically requiring the prior estimation of rental and return distribution on historical data [37, 45, 50, 97]. Inventory intervals have been computed in a similar fashion [53, 106], selecting those that minimize the likelihood of a station becoming either empty or full. Finally, different to those approaches, [21] and [54] simulate the performance of several sets of initial inventories and select the one that performs best.

All works cited above compute target inventories and inventory intervals based on historical trip demand. However, they disregard external factors such as weather conditions, the importance of which has been widely acknowledged in the literature (see, e.g., [28, 36, 39, 53, 61, 72, 87]). To this end, [53] extend the notion of service-levels proposed by [106], computing both target inventories and inventory intervals based on the service level and the predicted demand. The latter is estimated based on machine learning models trained on both temporal and weather data, therefore holding the potential of providing better performing target inventories and inventory intervals. The authors also introduce two additional hyperparameters, α and β, allowing operators to align inventory intervals and target inventories with their priorities for either rentals or returns (see Appendix B for more details), hence making it an attractive approach to operators.

5.2.3 Demand prediction for BSSs

Demand prediction is an essential step in the rebalancing process, enabling operators to anticipate which stations require higher inventories to better serve trip demand. Accurately predicting rentals and returns is challenging, as it is influenced by numerous factors [119]. Literature on demand prediction in BSSs can be divided into approaches predicting at global demand level (see e.g., [39, 104, 127]), at the level of station-clusters [**?**] or city regions(see e.g., [5, 17, 31, 119]), and at the level of individual stations (see e.g., [6, 25, 53, 89, 101]). We here focus on works predicting demand at station level, which is required for rebalancing operations since they are tailored considering the rentals and returns for each station.

Different techniques have been used to predict demand in BSSs. The average of historical trips (i.e., rentals and returns) can be used to estimate future demand [2, 21, 41, 45, 97, 106, 129],

which can be seen as a naive predictor. Alternatively, rental and return (often Poisson) distributions have been estimated from historical trip data, from which trip demand is then sampled and used within the optimization models (see, e.g., [43, 73, 129, 131]). While such prediction methods centered on historical mean demand have been quite popular, several studies (see, e.g., [34, 72, 74] haven suggested that such methods may result in high prediction errors when compared to ML models, as the latter can take into considerations features beyond historical trip data, including weather conditions, time of day, day of the week and the occurrence of special events.

Generally, ML algorithms can capture intricate patterns and correlations and may, therefore, result in significantly more accurate demand predictions. Random forests and gradient boosted trees, in particular, have been found to provide competitive prediction accuracy in the context of rental and return predictions [6, 53, 74, 125, 127]. Typically, such models integrate weather and temporal features, highlighting their importance to accurately predict demand. Here, [53] utilize Singular Value Decomposition, a dimension reduction technique, to reduce the dimensionality of the trip data, eliminate noise and improve both time required and the accuracy of hourly station-level rental and return predictions (details can be found in Appendix B). These reasons make this model an attractive option to estimate the future trip demand for our rebalancing optimization models.

5.2.4 Reoptimization modes: static, rolling, and folding

In practice, both single-period and multi-period models can be implemented in several ways. A simple, yet common approach (see, e.g., [41, 62, 75, 105, 109, 130]) is to optimize once over the entire planning horizon (i.e., all considered time-periods) and then implement all rebalancing decisions (i.e., the number of bikes dropped off and picked up at each station) as planned for all time periods.

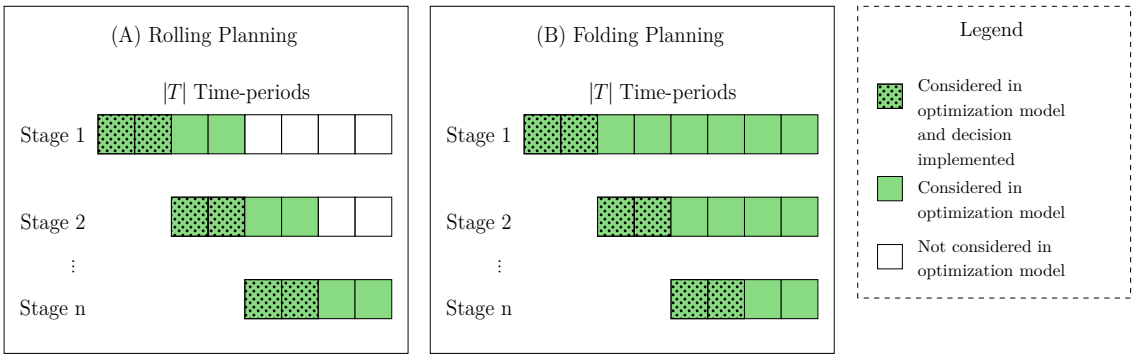

Figure 5.1 The structure of rolling and folding planing

While multi-period models can represent the consequences of decisions made at early time-periods, when executed within a static planning, they do not benefit from updated system information (such as station inventories or improved demand predictions). Therefore, practitioners often tend to reoptimize the rebalancing decisions throughout the planning horizon. Specifically, rolling planning (also called rolling window planning) considers the reoptimization over several time-periods at predefined reoptimization stages. Figure 5.1 (A) depicts the rolling planning, where, at each reoptimization stage, the green squares indicate time periods considered in the optimization models, and the green dotted squares indicate decisions of the time-periods actually executed in practice. This approach has been popular, with works implementing rolling planning for both multi-period (see, e.g., [44, 73, 82, 101, 110, 129]) and single-period models (see e.g., [42, 45, 49]). This allows for correcting ineffective planning, e.g., due to forecasting inaccuracies. In addition, considering only a subset of all time-periods within the planning window has the advantage of resulting in a more tractable optimization model.

If model decisions at early time-periods may impact decisions several time-periods ahead, one may want to consider a model with a longer planning horizon. Folding planning therefore optimizes on the remaining planning horizon, as depicted in Figure 5.1 (B). While this is generally a common approach in multi-period models [71], to the best of our knowledge, this approach has not yet been explored for BSSs rebalancing planning.

5.3 Dynamic Rebalancing Models

In this section, we formulate the DBRP as MIP models. Section 5.3.1 first describes a multi-period model, which minimizes the unmet rental and return demand, the more popular objective in the literature (as discussed in Section 5.2.1). We then propose two models integrating inventory intervals and target inventories into the objective functions in Section 5.3.2.

5.3.1 Multi-period rebalancing model minimizing lost demand

We consider the model with multiple time-periods from [70]. The input parameters are listed in Table 5.1 and decision variables are shown in Table 5.2. We denote S as the set of stations, while V denotes the set of available vehicles. Each station $s \in S$ has a capacity of C_s docks and each vehicle $v \in V$ can hold at most $\hat{C}_v$ bikes. We consider a planning horizon with $|T|$ time-periods, where each time period $t \in T$ represents a duration of L_t minutes.

We assume that a vehicle can visit only one station per time-period. As a result, a vehicle can visit at most T stations during the entire planning horizon. The decision variables $r^{+,t}_{s,v}$ and

Table 5.1 Input parameters of the optimization model

Input Parameters	Definition
S	The set of stations.
V	The set of vehicles.
T	The set of discreted time-periods.
C_s	The capacity of station $s \in S$.
$\hat{C}_v$	The capacity of vehicle $v \in V$.
L_t	The duration (in minutes) of time period $t \in T$.
d_s^1	The initial number of bikes at the station $s \in S$.
$\hat{d}_v^1$	The initial number of bikes in vehicle $v \in V$.
$z_{s,v}^1$	The initial location of each vehicle $v \in V, s \in S$.
$f_s^{+,t}$	The expected rental demand at station $s \in S$ in period $t \in T$.
$f_s^{-,t}$	The expected return demand at station $s \in S$ in period $t \in T$.

Table 5.2 Decision variables of the optimization model

Variables	Definition
d_s^t	The number of bikes available at station $s \in S$ at the beginning of period $t \in T$.
$\hat{d}_v^t$	The number of bikes in vehicle $v \in V$ at the beginning of period $t \in T$.
$x_s^{+,t}$	The number of successful rentals starting from station $s \in S$ in period $t \in T$.
$x_s^{-,t}$	The number of successful returns ending at station $s \in S$ in period $t \in T$.
$r_{s,v}^{+,t}$	The number of bikes picked up at station $s \in S$ by vehicle $v \in V$ in period $t \in T$.
$r_{s,v}^{-,t}$	The number of bikes dropped off at station $s \in S$ by vehicle $v \in V$ in period $t \in T$.
$z_{s,v}^t$	$z_{s,v}^t = 1$, if vehicle $v \in V$ visits station $s \in S$ in period $t \in T$; 0 otherwise.

$r_{s,v}^{-,t}$ represent the number of bikes vehicle v picks up and drops off, respectively, at station s during period t. Furthermore, binary variable $z_{s,v}^{t}$ takes value 1 if and only if vehicle v visits station s at time-period t. For each time-period, intermediate variables are used: the number of bikes available at stations and in vehicles, successful trips, and vehicle routes. The resulting MIP model is expressed as follows:

$$\min \quad \sum_{s\in S}\sum_{t\in T}(f_s^{+,t}-x_s^{+,t})+\sum_{s\in S}\sum_{t\in T}(f_s^{-,t}-x_s^{-,t}) \tag{5.1}$$

$$\text{s.t.} \quad \hat{d}_v^{t+1}=\hat{d}_v^{t}+\sum_{s\in S}(r_{s,v}^{+,t}-r_{s,v}^{-,t}) \qquad \forall\ v\in V,\ t\in T \tag{5.2}$$

$$d_s^{t+1}=d_s^{t}-\sum_{v\in V}(r_{s,v}^{+,t}-r_{s,v}^{-,t})-x_s^{+,t}+x_s^{-,t} \qquad \forall\ s\in S,\ t\in T \tag{5.3}$$

$$\sum_{s\in S} z_{s,v}^{t}=1 \qquad \forall\ v\in V,\ t\in T \tag{5.4}$$

$$r_{s,v}^{+,t}+r_{s,v}^{-,t}\le \hat{C}_v z_{s,v}^{t} \qquad \forall\ s\in S,\ v\in V,\ t\in T \tag{5.5}$$

$$0\le \hat{d}_v^{t}\le \hat{C}_v \qquad \forall\ v\in V \tag{5.6}$$

$$0\le d_s^{t}\le C_s \qquad \forall\ s\in S \tag{5.7}$$

$$0\le x_s^{+,t}\le f_s^{+,t}, 0\le x_s^{-,t}\le f_s^{-,t} \qquad \forall\ s\in S,\ t\in T \tag{5.8}$$

$$0\le r_{s,v}^{+,t}, r_{s,v}^{-,t}\le \hat{C}_v \qquad \forall\ s\in S,\ v\in V,\ t\in T \tag{5.9}$$

$$z_{s,v}^{t}\in\{0,1\} \qquad \forall\ s\in S,\ v\in V,\ t\in T. \tag{5.10}$$

The objective function (5.1) minimizes the total lost demand, i.e., the unmet expected demand for both rentals and returns, over the entire planning horizon at all stations. Constraints (5.2) compute the number of bikes in each vehicle v in period $t+1$ based on the number of bikes in the previous period and the number of picked up/ dropped off bikes. Constraints (5.3) compute the number of bikes in period $t+1$ at each station s as the sum of the number of bikes of that station in the previous period, the number of bikes rebalanced by vehicles, and those moved by users (i.e., successful rentals and returns). Constraints (5.4) ensure that each vehicle v can only be at one station at each time-period. Constraints (5.5) ensure that a vehicle can perform operations at a station only when it is present at that station. Constraints (5.6) impose that the number of bikes in each vehicle is bounded by its capacity. Constraints (5.7) are the capacity constraints for the stations. Constraints (5.8) bound the number of successful trips by the expected rental and return. Finally, constraints (5.9) enforce that the pick-up and drop-off operations respect the vehicle's capacities.

The above model, denoted as DROB-LD, derives rebalancing strategies for the entire planning horizon, i.e., it decides how many bikes each vehicle should pick up or drop off at which station. The model can be easily implemented in different reoptimization modes (static, rolling, and folding planning) with alterable length of planning horizons and duration of time-periods, depending on the requirements of the decision-maker.

5.3.2 Rebalancing models based on inventory interval and target inventory

Even though the above used objective minimizing the lost demand is quite popular in the literature, its performance is sensitive to the accuracy of the expected rentals $f_s^{+,t}$ and returns $f_s^{-,t}$. Rather than minimizing the deviation from such a point estimate, we propose to minimize the deviation from either the inventory interval or the target inventory. This approach provides a buffer for the station inventories, allowing them to maintain reasonable inventories even when the trip prediction is less accurate, and to be better prepared for fluctuations of the stochastic demand. To this end, we propose two multi-period models with novel objective functions: **D**ynamic **R**ebalancing **O**ptimization for **B**SS based on **T**arget Inventories (DROB-T) and DROB-I. The parameters and variables used in both models are depicted in Table 5.3.

Table 5.3 Parameters and variables that define inventory intervals and target inventories

Input	Definition
ℓ_s^t	The target inventory of station $s \in S$ at period $t \in T$
$\underline{\ell}_s^t$	The lower bound for the inventory interval of station $s \in S$ at period $t \in T$
$\overline{\ell}_s^t$	The upper bound for inventory interval of station $s \in S$ at period $t \in T$
Variables	**Definition**
$e_s^{+,t}$	The number of bikes above the upper bound at station $s \in S$ in period $t \in T$
$e_s^{-,t}$	The number of bikes below the lower bound at station $s \in S$ in period $t \in T$

The objective function of DROB-T (5.11) aims at minimizing the total deviations between station inventories and target values, thus yielding the following formulation:

$$\begin{aligned} \min \quad & \sum_{s\in S}\sum_{t\in T} |\ell_s^t - d_s^t| && (5.11)\\ \text{s.t.} \quad & (5.2) - (5.10). \end{aligned}$$

DROB-I is designed to keep the station inventories as much as possible within the computed intervals. To this end, DROB-I is formulated as the objective function (5.12), along with constraints (5.2)-(5.10) and (5.13)-(5.15) as defined below:

$$\begin{aligned} \min \quad & \sum_{s\in S}\sum_{t\in T} e_s^{-,t} + e_s^{+,t} && && (5.12)\\ \text{s.t.} \quad & \underline{\ell}_s^t - e_s^{-,t} \le d_s^t && \forall\ s \in S,\ t \in T && (5.13)\\ & d_s^t \le \overline{\ell}_s^t + e_s^{+,t} && \forall\ s \in S,\ t \in T && (5.14)\\ & e_s^{-,t} \ge 0, e_s^{+,t} \ge 0 && \forall\ s \in S,\ t \in T && (5.15)\\ & (5.2) - (5.10). \end{aligned}$$

Figure 5.2 exemplifies the inventory of a station s, as well as the lower bound and upper bound of the inventory interval, and its target inventory. For DROB-I, the excess of inven-

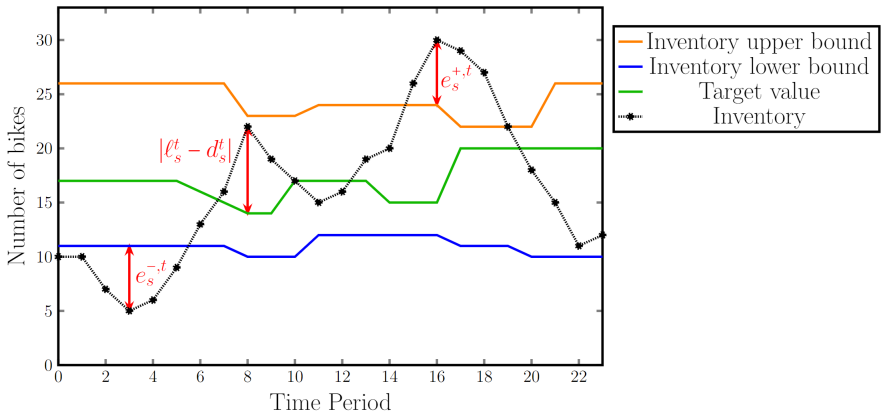

Figure 5.2 Example of deviations from inventory intervals and target inventory

tory at each time-period ($e_s^{+,t}$ and $e_s^{-,t}$, respectively), with respect to the inventory interval $[\underline{\ell}_s^t, \overline{\ell}_s^t]$, is computed by constraints (5.13) and (5.14). For DROB-T, Figure 5.2 also illustrates the deviation $|\ell_s^t - d_s^t|$ from the current target value. By minimizing these deviations, DROB-I and DROB-T aim at providing safety buffers to the station inventory and are therefore more likely of being capable to deal with stochastic demand fluctuations. We next present two toy examples to provide an intuition of the potential benefits of the two new objective functions.

Table 5.4 Comparative analysis of three objective functions for rebalancing operations

Objective Function	Example 1		Example 2	
	Dropped off Bikes	Expected Lost Demand*	Dropped off Bikes	Expected Lost Demand*
DROB-LD (5.1)	1	0.33	0	0.67
DROB-T (5.11)	8	0	5	0
DROB-I (5.12)	2	0	2	0

* The expected lost demand is calculated considering all possible chronological sequences of rentals and returns derived from historical trip data. For simplicity, we assume that the probability of a rental occurring before a return equals to that of a return happening before a rental.

Example 1. Consider an empty station with 10 docks. For a given time-period at a given day, historical rentals follow a uniform distribution ranging from 0 to 2, while no returns have been observed. A predictive model is used to predict the expected demand and target value and inventory interval are computed as to ensure a sufficient service level (see Section 5.4.1 and Appendix B for details). As a result, an estimation of 1 rental and no return is obtained, directly used in model DROB-LD. For DROB-T, the computed target value is 8, while for

DROB-I, the computed inventory interval is [2, 10]. Table 5.4 summarizes the number of bikes dropped off at that station according to each of the three models. Furthermore, the table reports the expected lost demand if rental demand is uniformly distributed between 0 to 2. Here, DROB-T and DROB-I drop off at least 2 bikes, accounting for the potential demand of 2 rentals. In contrast, DROB-LD drops off only 1 bike, and therefore lacks 1 bike when the rental demand is 2.

Example 2. Consider that the same empty station, for another time-period and given day, has a uniform distribution between 0 and 2 for both rentals and returns. Both the estimated rental and return are therefore 1. The computed target value is 5, whereas the inventory interval is [2, 8]. In this case, DROB-T and DROB-I still drop off at least 2 bikes and therefore do not induce any unmet rental demand. In contrast, DROB-LD implicitly assumes that returns cancel rentals, and thus does not drop off any bikes, which may result in unmet demand when the rental demand is 1 or higher.

5.4 Experiments and Results

We now employ computational experiments to explore the benefits of the proposed models. Section 5.4.1 introduces the synthetic data and reports on the corresponding empirical results. A case study on real-world data is then presented in Section 5.4.2.

Computational environment. All optimization models are solved using IBM ILOG CPLEX v20.1.0.0 on 2.70 GHz Intel Xeon Gold 6258R machines with 8 cores. Optimization terminates once the MIP gap reaches 0.01% or the time limit of 24 hours is reached.

5.4.1 Experiments on synthetic data

We here focus on experiments carried out synthetic problem instances. To this end, Section 5.4.1 first introduces the instance generator that generates weather-dependent trip data. Section 5.4.1 then details the experimental set-up, including the machine learning model used to predict rental demand, the computation of inventory intervals and targets, and the simulator. This section also summarizes the computational results, comparing the performance of the various planning models. Finally, Section 5.4.1 then explores the performance of the planning models under the assumptions that predictions are less accurate.

Synthetic dataset

Even though we have access to real-world trip data, we synthetically generate instances for several reasons. First, the available real-world trip data lacks information on unobserved demand. Second, existing data may contain noise related to trip and station inventory data. Finally, rebalancing operations conducted by operators impact station inventory, but data on such operations is not openly available.

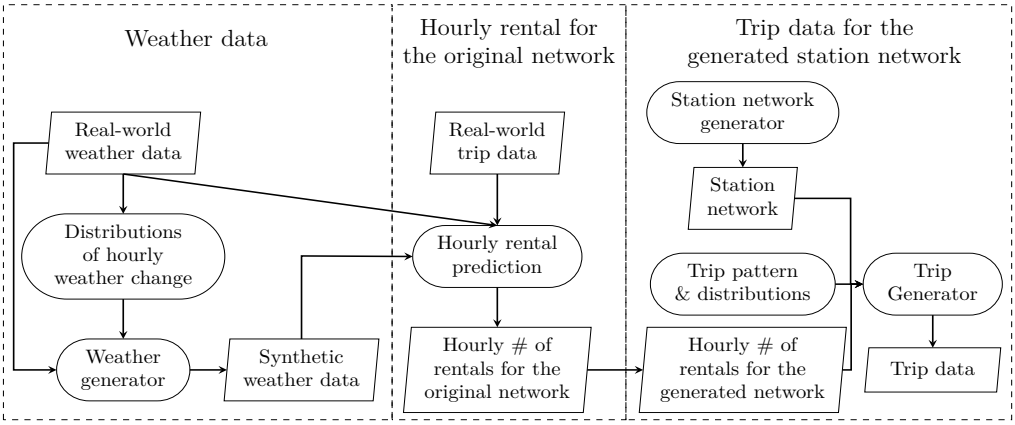

Figure 5.3 Process to generate weather data, station network and trip data

The general data generation process of the here employed instance generator is depicted in Figure 5.3. We extend the instance generator proposed by [70] (corresponding to the third box in Figure 5.3: "*Trip data for the generated station network*"), which is capable of generating diverse station networks and trip data, based on predefined *trip patterns & distributions* that align with those observed in real-world BSSs. Whereas the instance generator of [70] generates trip data under the same weather conditions (specifically, assuming high demand during summer months), we here explicitly acknowledge the strong correlation between weather conditions and trip demand. As such, we extend this instance generator as follows. (i) We introduce varying weather conditions into our generator (corresponding to the first box "*Weather data*" in Figure 5.3) by estimating statistical distributions that represent the hourly changes of weather conditions estimated on real-world weather data, and then sample new weather conditions from these distributions. (ii) We then compute the system-wide level of trip demand that is correlated to such weather conditions (corresponding to the second box "*Hourly rental for the original network*" in Figure 5.3), i.e., the hourly number of rentals for the entire network. Finally, individual trip demand is generated for the generated station network. Each of these components is next explained in detail.

Generation of weather data. Given that we ultimately aim at generating trip data that resembles periods of the demand peak season, we are interested in generating weather conditions for the months of June to August, which tend to have high user demand. We therefore procure *Real-world weather data*[1] from Montreal for June, July, and August from 2017 to 2020, resulting in a total of 368 days (including both weekdays and weekends). We select two features of utmost importance (see, e.g., [28, 36, 39, 61]): temperature and humidity. We analyze the temperature and humidity differences between consecutive hours throughout the day and divide the day into four distinct time segments such that temperature and humidity tend to remain relatively stable within each of which (0 am – 5 am, 6 am – 11 am, 12 pm – 5 pm, and 6 pm – 11 pm). The hourly differences are then used to estimate normal *Distributions of hourly weather change* for each time segment. The estimated distributions of temperature change for each time segment can be found in Appendix C.

Synthetic weather data has been generated for a total of 500 days as follows. For each hour of the original 368 days, we use the temperate and humidity that originally occurred at that day and add a change of temperature and humidity, respectively, sampled from the corresponding distributions. To obtain a total of 500 days, we repeat the process with the first 132 of the original 368 days. To ensure that the generated weather conditions are sufficiently realistic and avoid drastic fluctuations, we introduce constraints to keep the temperature and humidity within 5.5°C–36°C and 15%–99%, respectively. We denote the final set of the 500 generated days with synthetic weather data as the *Synthetic weather data*.

Generation of hourly rental demand. Using the temperature, humidity, hour, and weekday as features, we estimate the *Hourly rental prediction* model. This linear regression model is trained using the *Real-world weather data* and *Real-world trip data* (also see Section 5.4.2) and captures the correlation between time and weather conditions and the total demand level for the entire station network. The trained regression model is then used to estimate the total system-wide number of *Hourly # of rentals for the original network* that depends on the *Synthetic weather data* throughout each day. Given that this total number of rentals has been estimated on the original network from the *Real-world trip data* (here, the BIXI network with over 600 stations), this number is then scaled to the number of stations used in the here considered *Station network* (which has 60 stations). These hourly system-wide rental demands (*Hourly # of rentals for the generated network*) then serve as input to generate the detailed trip data.

[1] https://climate.weather.gc.ca

Generation of station network and individual trip data. We generate two ground truth problem instances, denoted GT1 and GT2, each of which contains a *station network*, as well as hourly weather data and detailed trip information for 500 days. In both instances, the network contains 60 stations with different numbers of city center stations. The stations within city centers are equipped with 40 docks, while those outside city centers have 20 docks each.

Generated trips contain the origin station, the destination station, the departure time, and the arrival time. We consider four trip patterns of user behaviors with origins and destinations outside (O) and inside (I) city centers: (i) users who live outside city centers and work inside city centers typically use similar origin (outside city centers) and destination stations (inside city centers) during peak hours (OI trips); (ii) users who live and work outside city centers (OO trips); (iii) random non-work related trips occurring during the day (RD trips), and (iv) random non-work related trips occurring during the night (RN trips).

Table 5.5 Characteristics of the two considered ground truth instances

Instance		GT1	GT2
Network (60 stations)	# of city centers	1	2
	# of stations per city center	9	6
	City center capacity	26%	35%
Trip Pattern	OI	32%	32%
	OO	32%	32%
	RD	23%	23%
	RN	13%	13%

The characteristics of instance ground truths GT1 and GT2 are described in Table 5.5. Although the proportions of work-related (i.e., city center related) trips are identical in GT1 and GT2, the latter has more city center stations (12 stations in 2 city centers, as opposed to 9 stations in 1 city center). As such, work related trips in GT2 are distributed over a larger number of stations, which are therefore less stressed.

To sample individual trips for the considered *station network*, we assume that each trip type follows a particular temporal distribution, indicating the probabilistic time at which the rental occurs (as detailed in [70]). For each day, the *trip generator* then sequentially samples trips (i.e., origin-destination pairs and exact time stamps) from the *Trip pattern & distributions* (as defined by the ground truth) until the *hourly # of rentals for the generated network* is met, resulting in the final set of *Trip Data*.

Model performance based on regular trip predictions

Experimental set-up. The 500 generated days for GT1 and GT2 are separated into training set, validation set and test set as follows. The first 250 days are allocated to calibrate the gradient boosted tree introduced in [53], capable of predicting hourly station demand and trained on trips (time of rental and arrival/departure stations), weather conditions (temperature and humidity), and temporal data (day of the week, hour of the day, and a binary indicator for holidays). The subsequent 100 days are used for the validation and fine-tuning of the gradient boosted tree and inventory intervals. Details on the training of the gradient boosted tree can be found in Appendix B.

The remaining 150 days constitute the test set on which the optimization algorithms are executed. We consider a planning horizon from 7 a.m. to 3 p.m., discretized into 8 time-periods, each with a duration of one hour. For each of the 150 days, we assume to have access only to its corresponding weather conditions, but not to the exact trip (i.e., rental and return) demand. This is a reasonable assumption in practice, where one can assume to have access to a reasonably accurate weather prediction. Based on such weather conditions, the trained gradient boosted tree then predicts the hourly rental and return demand for each station ($f_s^{+,t}$ and $f_s^{-,t}$), used within model DROB-LD. The inventory intervals and target values, used within models DROB-I and DROB-T, are then computed based on the predicted rental demand (details can also be found in Appendix B). For each of the 150 days, the rebalancing planning solutions provided by the various models are then evaluated in the simulator (see Section 5.4.1) on the exact trip data. Note, again, that the optimization models only have access to demand predictions (based on weather data), whereas the simulator evaluates on the exact trip demand of the days in the test set.

In all experiments, 4 vehicles are available to rebalance the stations, each with a capacity for 40 bikes. The initial inventory of stations is obtained by solving an overnight rebalancing problem (equivalent to the one used in [70]).

Each of the optimization models can be executed in different reoptimization planning modes. In *static planning*, the optimization model is solved once for the entire planning horizon. The rebalancing strategies of the first 6 (out of 8) time-periods are then executed within the simulator to estimate the lost demand. The *rolling planning* has 3 optimization stages, each of which contemplates 4 time-periods. At each stage, the rebalancing decisions of the first 2 time-periods are executed within the simulator, as illustrated in Figure 5.1 (A). The *folding planning* uses all the remaining time-periods at each stage, as depicted in Figure 5.1 (B). The fine-grained discrete-event simulator from [70] here used employs a chronological first-arrive-first-serve rule, for both user rentals and returns, as well as rebalancing vehicles (i.e.,

pick-ups and drop-offs). Events are discretized events into 1-minute time-slots, which results in a particularly detailed and realistic simulation.

Computational Results. Table 5.6 illustrates the average lost demand and computing time (over the test set) for all the models and reoptimization modes (**S**tatic, **R**olling, and **F**olding). We report the computing times required to solve the optimization models as 'Opt. Time' (in minutes). The lost rental demand is computed as the relative gap between successful rentals and the original rental demand specified in the instances over the entire planning horizon, i.e., $\frac{\sum_{s,t}(f_s^{+,t}-\hat{x}_s^{+,t})}{\sum_{s,t} f_s^{+,t}}$, where $\hat{x}_s^{+,t}$ is the number of successful rentals in the simulator. The lost return demand is computed as $\frac{\sum_{s,t}(\hat{x}_s^{+,t}-\hat{x}_s^{-,t})}{\sum_{s,t} \hat{x}_s^{+,t}}$, where $\hat{x}_s^{-,t}$ is the number of successful returns in simulator. Since, in practice, return demand does not exist when the corresponding rental demand is unsuccessful, the lost returns are only associated with successful rentals $\hat{x}_s^{+,t}$. We also present the relative difference ($\Delta(\%)$) of the rental, return, and total lost demand of DROB-I and DROB-T when compared to DROB-LD under the respective reoptimization mode.

Table 5.6 Results of dynamic rebalancing models for GT1 and GT2 under regular prediction

Instance	Model	Reopt. Mode	Opt.Time (min.)	Lost Demand (%)					
				Rental	Δ(%)	Return	Δ(%)	Total	Δ(%)
GT1	DROB-LD	*S*	0.19	13.14		2.74		8.31	
		R	0.02	11.94		2.76		7.65	
		F	0.21	11.54		2.87		7.48	
	DROB-I	*S*	0.18	10.68	▼ -18.72	3.79	▲ 38.32	7.43	▼ -10.59
		R	0.10	8.18	▼ -31.49	2.90	▲ 5.07	5.66	▼ -26.01
		F	0.39	8.39	▼ -27.30	2.84	▼ -1.05	5.74	▼ -23.26
	DROB-T	*S*	1.53	9.98	▼ -24.05	3.97	▲ 44.89	7.14	▼ -14.08
		R	0.15	8.20	▼ -31.32	1.52	▼ -44.93	5.01	▼ -34.51
		F	1.79	8.21	▼ -28.86	1.61	▼ -43.90	5.06	▼ -32.35
GT2	DROB-LD	*S*	0.41	8.27		1.68		5.13	
		R	0.06	8.16		1.67		5.06	
		F	0.44	8.11		1.62		5.01	
	DROB-I	*S*	0.04	7.76	▼ -6.17	2.51	▲ 49.40	5.24	▲ 2.14
		R	25.24	6.31	▼ -22.67	1.35	▼ -19.16	3.92	▼ -22.53
		F	229.15	5.87	▼ -27.62	1.12	▼ -30.86	3.57	▼ -28.74
	DROB-T	*S*	0.13	8.79	▲ 6.29	0.63	▼ -62.50	4.91	▼ -4.29
		R	0.28	6.96	▼ -14.71	0.58	▼ -65.27	3.89	▼ -23.12
		F	0.32	6.97	▼ -14.06	0.57	▼ -64.81	3.90	▼ -22.16

Table 5.6 allows for the following observations:

1. **Comparison of proposed models.** From Table 5.6, models DROB-I and DROB-T generally outperform DROB-LD for both GT1 and GT2. For example, on GT1, DROB-I reduces the lost demand from 7.65% to 5.66% in rolling planning, while being solved within seconds. While under static planning, DROB-I and DROB-T tend to outperform DROB-LD, they consistently outperform DROB-LD by a higher rate under rolling planning (reducing total lost demand by 23.26%-34.51%)

2. **Comparison of different reoptimization modes.** The rolling and folding planning consistently result in lower lost demand than the static planning, likely due to the fact that they update the station inventories before reoptimizing at every reoptimization stage (see Figure 5.1). Updating the inventory narrows the gap between the estimated inventories in the optimization model and the observed inventories during the simulation, allowing the optimization models to make more informed decisions. For example, rolling and folding planning in DROB-I on GT1 reduce the lost demand from 7.43% to 5.66% and 5.74%, respectively, over the static planning. Note that updating weather forecast may have significat impact in the rolling planning, which will be discussed in Section 5.4.1. In terms of computing times, even though most models have been solved within 1-2 minutes, the folding planning requires much longer computing times for GT2.

3. **Comparison between GT1 and GT2.** While the general conclusions and tendencies are the same for GT1 and GT2, for GT1, models present more unmet demand than for GT2. Indeed, GT2 has more stations located in city centers, a region with high work-related demand, resulting in more evenly distributed trip patterns for these stations. Moreover, the larger number of city center stations translates into greater availability of docks in the network, as stations located in this area contain twice the number of docks than stations located in other regions of the city.

Results based on perfect trip prediction. We further carry out experiments with perfect trip information for each day of the test set, i.e., the model optimizes on the exact rental and return demand on which its rebalancing policy is later evaluated within the simulator. These experiments establish an empirical performance bound and provide insights into the efficiency of the optimization models and reoptimization modes. Figure 5.4 presents the total lost demand obtained using the predicted trip demand (i.e., the same as used for Table 5.6) and the perfect information for GT1. Unsurprisingly, if perfect information was available, DROB-LD would consistently outperform DROB-I and DROB-T, given that an inventory safety buffer would be unnecessary. However, in reality, demand is stochastic. In this case, DROB-I and DROB-T can provide more robust station inventories, allowing them to deal

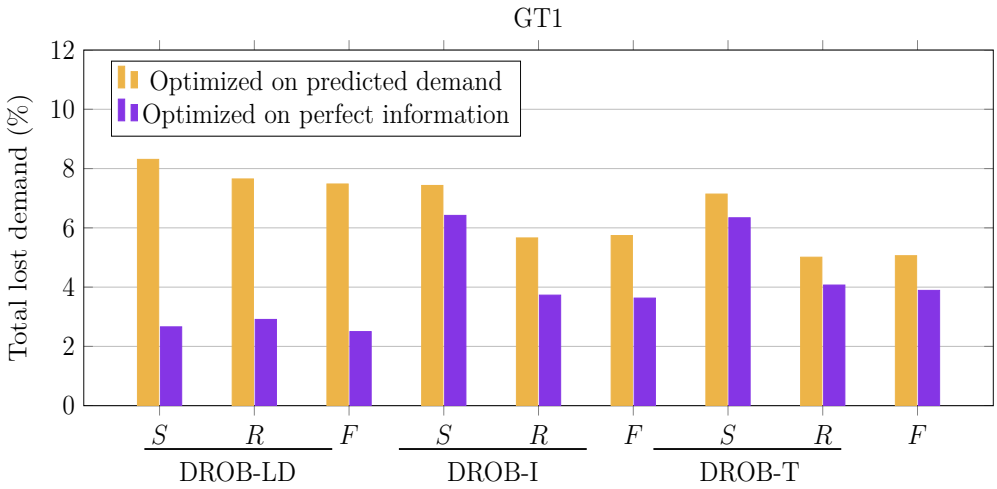

Figure 5.4 Results of total lost demand for GT1 using regular prediction and ground truth

with the stochastic trip demand. Interestingly, DROB-I and DROB-T still benefit from inventory updates (rolling and folding planning) under perfect information, as opposed to DROB-LD. Indeed, the former two models rely on the current inventory levels to update their objective functions, while the objective of DROB-LD remains unchanged, even when station inventories change. As a result, reoptimization for DROB-LD is not beneficial. Finally, the results also enable us to derive insights into the empirical bounds on the potential gains achieved through the utilization of a more accurate predictive model. While the gains are substantial (~5% of lost demand) for DROB-LD, they are much smaller (~1-2%) for DROB-I and DROB-T. While using a more accurate predictions may obviously lead to reduced unmet demand, we will next investigate how those models perform when predictions are less accurate.

Model performance based on noisy prediction

The optimization models used in our previous experiments have taken as input demand predictions and interval predictions that have assumed a perfect weather forecast. In practice, weather forecasts for the next 2 to 8 hours can be prone to inaccuracies. In a similar vein, having access to a predictive model with sufficiently high accuracy may not always be possible. We will now investigate the performance of the various models under the assumption that demand and interval predictions are less accurate. To this end, we deliberately introduce noise into the performed trip predictions. Since accurately predicting demand becomes increasingly challenging as we project further into the future, we introduce more noise to later time-periods.

Noisy predictions. We consider two types of effects caused by noises over demand predictions: (i) overestimation, e.g., due to a forecast of overly favorable weather conditions and, therefore, expecting a higher number of trips than will actually occur; (ii) underestimation, e.g., due to a forecast of adverse weather conditions, therefore predicting a lower number of trips than will actually occur.

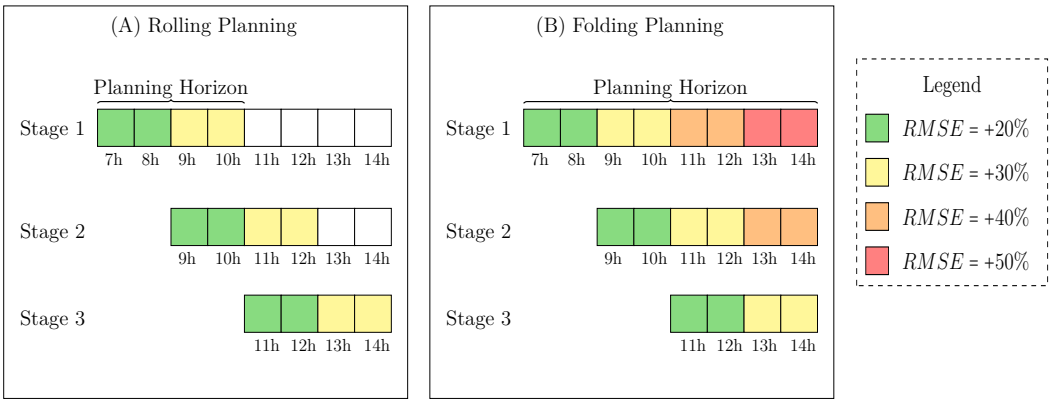

Figure 5.5 Noise added to the demand prediction over the planning horizon of the rolling and folding planning

To create noisy predictions for the demand at station s at time t, we sample noise from a normal distribution. Its mean μ is determined by the original predicted number of rentals (or returns), while its standard deviation is strategically adjusted to achieve a predetermined increase in the Root Mean Squared Error (RMSE) for these noisy predictions in comparison to μ. This approach is applied throughout the planning horizon for each stage in the rolling and folding planning as illustrated in Figure 5.5. Values sampled above the mean are used to create overestimating predictions, whereas values below the mean are sampled to create underestimating predictions. Thus, an underestimating forecast consistently predicts lower demand and an overestimating forecast consistently predicts higher demand. Note that static planning operates under the same conditions as the first stage of folding planning.

Results. Tables 5.7 and 5.8 show the results for the three optimization models under underestimating and overestimating predictions, respectively. Same as in Table 5.6, the relative difference of DROB-I and DROB-T when compared to DROB-LD is reported ($\Delta(\%)$).

Based on Tables 5.7 and 5.8, we summarize our observations as follows:

1. **Comparison of models.** Although DROB-LD shows performance improvement in the case of overestimating predictions, DROB-I and DROB-T consistently demonstrate

Table 5.7 Results of dynamic rebalancing models for GT1 and GT2 under underestimating predictions

Instance	Model	Reopt. Mode	Opt.Time (min.)	Lost Demand (%) Rental	Δ(%)	Return	Δ(%)	Total	Δ(%)
GT1	DROB-LD	S	0.02	13.97		4.56		9.62	
		R	0.02	11.90		3.54		7.99	
		F	0.04	11.57		3.56		7.81	
	DROB-I	S	11.45	10.31	▼ -26.20	4.37	▼ -4.17	7.51	▼ -21.93
		R	0.13	8.23	▼ -30.84	2.82	▼ -20.34	5.65	▼ -29.29
		F	21.13	7.73	▼ -33.19	1.68	▼ -52.81	4.84	▼ -38.03
	DROB-T	S	0.20	11.89	▼ -14.89	2.55	▼ -44.08	7.52	▼ -21.83
		R	0.30	9.98	▼ -16.13	1.11	▼ -68.64	5.78	▼ -27.66
		F	2.84	10.03	▼ -13.31	1.04	▼ -70.79	5.78	▼ -25.99
GT2	DROB-LD	S	0.01	11.66		2.69		7.46	
		R	0.01	10.85		1.36		6.39	
		F	0.02	10.44		1.38		6.17	
	DROB-I	S	0.23	8.6	▼ -26.24	2.48	▼ -7.81	5.68	▼ -23.86
		R	0.19	7.02	▼ -35.30	1.40	▲2.94	4.32	▼ -32.39
		F	1.94	6.99	▼ -33.05	1.39	▲0.72	4.30	▼ -30.31
	DROB-T	S	21.30	9.42	▼ -19.21	1.27	▼ -52.79	5.55	▼ -25.60
		R	0.79	7.42	▼ -31.61	0.56	▼ -58.82	4.13	▼ -35.37
		F	21.88	7.48	▼ -28.35	1.16	▼ -15.94	4.45	▼ -27.88

lower lost demand in most cases. Especially within rolling and folding planning, DROB-I and DROB-T outperform DROB-LD considerably. Their advantage is particularly pronounced when optimizing on underestimating predictions.

2. **Comparison of different reoptimization modes.** Generally, the improvement of lost demand when transitioning from static planning to folding and rolling planning is more significant under perturbed trip predictions than under noise-free predictions (see Table 5.6). This confirms the importance of such reoptimization planning modes when less accurate predictions are used.

3. **Comparison between predictions.** Underestimating trip predictions results in higher lost demand compared to noise-free predictions, since fewer rebalancing operations are triggered. In contrast, overestimating predictions may lead to less lost demand, especially notable for DROB-LD. This is explained by the fact that overestimating predictions triggers more rebalancing operations in DROB-LD. We report the number of rebalancing operations (number of bikes picked up and dropped off) over

Table 5.8 Results of dynamic rebalancing models for GT1 and GT2 under overestimating predictions

Instance	Model	Reopt. Mode	Opt.Time (min.)	Lost Demand (%)					
				Rental	Δ(%)	Return	Δ(%)	Total	Δ(%)
GT1	DROB-LD	S	1.50	12.09		2.14		7.44	
		R	0.04	10.79		2.12		6.71	
		F	1.52	10.50		2.18		6.58	
	DROB-I	S	0.08	11.05	▼ -8.60	3.77	▲76.17	7.62	▲2.42
		R	0.06	8.75	▼ -18.91	1.62	▼ -23.58	5.35	▼ -20.27
		F	1.41	8.59	▼ -18.19	0.86	▼ -60.55	4.91	▼ -25.38
	DROB-T	S	0.06	11.05	▼ -8.60	3.94	▲84.11	7.71	▲3.63
		R	0.10	8.77	▼ -18.72	1.24	▼ -41.51	5.19	▼ -22.65
		F	0.26	8.75	▼ -16.67	1.21	▼ -44.50	5.16	▼ -21.58
GT2	DROB-LD	S	1.31	7.58		0.77		4.32	
		R	0.06	6.97		0.93		4.07	
		F	1.46	6.87		0.85		3.97	
	DROB-I	S	0.03	8.04	▲6.07	1.64	▲112.99	4.98	▲15.28
		R	0.03	7.00	▲0.43	0.60	▼ -35.48	3.93	▼ -3.44
		F	0.07	7.01	▲2.04	0.60	▼ -29.41	3.93	▼ -1.01
	DROB-T	S	0.08	8.63	▲13.85	1.56	▲102.6	5.27	▲21.99
		R	0.05	7.04	▲1.00	0.45	▼ -51.61	3.88	▼ -4.67
		F	0.16	7.06	▲2.77	0.46	▼ -45.88	3.89	▼ -2.02

the planning horizon in Figure 5.6. Indeed, overestimating predictions results in more rebalancing operations to meet the high demand. Overall, it appears that DROB-I and DROB-T are less sensitive to prediction noise than DROB-LD, given that they are designed to introduce a buffer into the optimized stations' inventories.

Remarks. Overall, the results indicate that DROB-LD is much more susceptible to the accuracy of demand predictions than DROB-I and DROB-T as the directly considers such predictions in its objective function. By introducing a buffer to the stations inventories, the performance of DROB-I and DROB-T remains more stable among the different prediction approaches. To visualize the improvement of rolling/folding planning over static planning, we illustrate the difference in total lost demand between static and rolling planning in Figure 5.7, and between static and folding planning in Figure 5.8 for both GT1 and GT2. It is worth noting that the difference in total lost demand in these figures consistently shows positive values, meaning that in all experiments, rolling and folding planning consistently outperform static planning. These improvements are even more significant in the experiments with underestimating and overestimating predictions. This can be attributed to the fact that, in

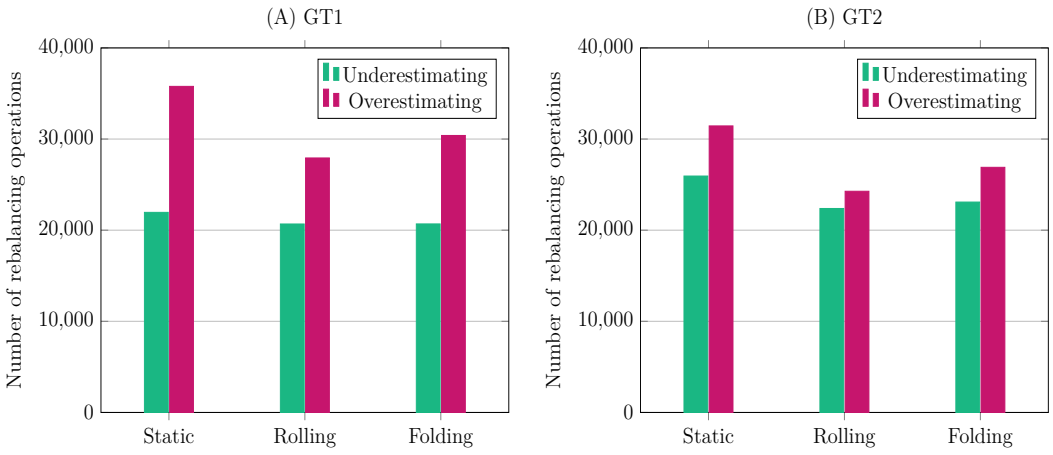

Figure 5.6 Number of rebalancing operations carried out in GT1 and GT2 for the DROB-LD

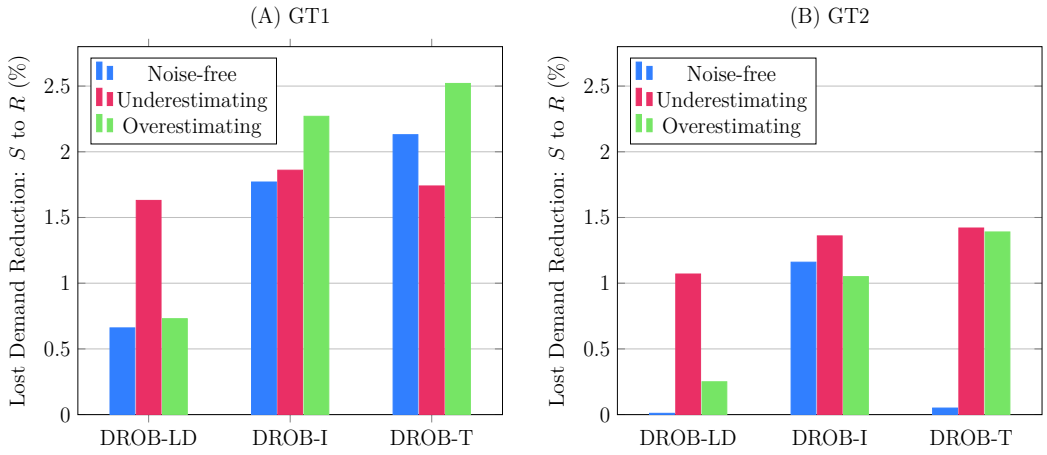

Figure 5.7 Improvement of the total lost demand from static to rolling planning

addition to updating the station inventory, a more accurate trip prediction is updated before re-optimization in each stage (see Figure 5.5).

5.4.2 Experiments on real-world data

In this section, we describe the real-world dataset used to validate the effectiveness of model DROB-I, which, on synthetic data, has demonstrated consistently low lost rentals while maintaining reasonable computing times. We first describe the real-world dataset in Section 5.4.2. We then describe the results in Section 5.4.2. The experimental set-up and planning horizon here considered are the same as in the experiments on synthetic data.

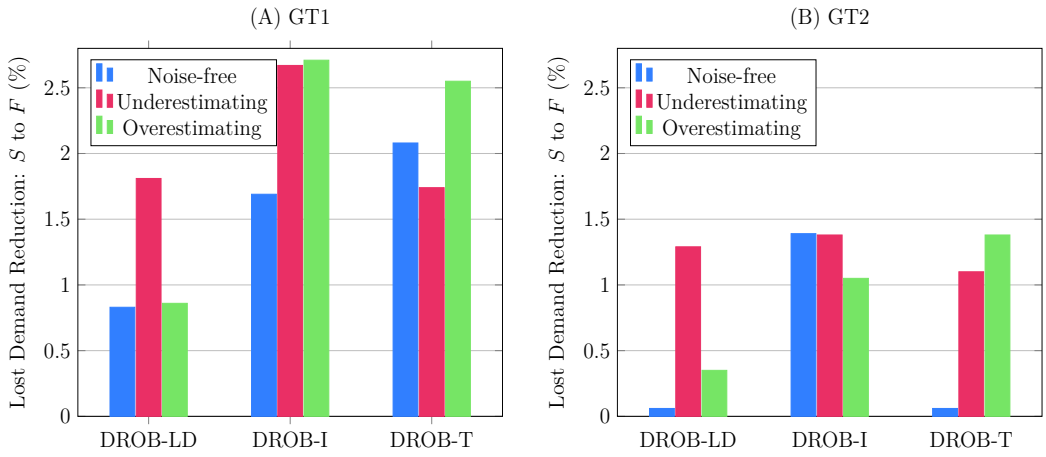

Figure 5.8 Improvement of the total lost demand from static to folding planning

Real-world dataset

The real-world data consists of weather, temporal, station, and trip data. Weather and temporal data are directly provided by the official website of the Government of Canada, including temperature and humidity. The temporal data contains the date, hour (0h -23h), year (2019), and weekday (Monday to Friday). The trip and station data are provided by BIXI[2]. The trip data contain the origin station, start time, destination station, and arrival time of each trip, while the station data contain the location and station capacity (i.e., the number of docks).

We only focus on trips during weekdays from May to September 2019. Selecting trips before 2020 ensures that analyzed trip patterns are not affected by the COVID-19 pandemic. Weekdays are chosen due to their typically consistent work-related trip patterns. The first 21 days of each month, excluding the weekends, constitute the training dataset. The remaining days of May are used for validation and the remaining days from June to September are assigned to the test dataset. The initial inventory for stations at 7 a.m. is also collected from BIXI dataset and serves as input for the optimization models.

For BIXI's station network, we exclude stations that have been relocated more than 1 km from their original locations by the operator during specific events, constructions, or holidays. As a result, 606 stations out of originally 620 remain in our experiments. This network is too large to be directly solved by general-purpose solvers. As a remedy, literature often divides the network into smaller clusters. Rebalancing is then performed within a specific cluster or between different clusters (see, e.g., [11, 32, 41, 51, 58, 72]).

[2]https://bixi.com/en/open-data-2/

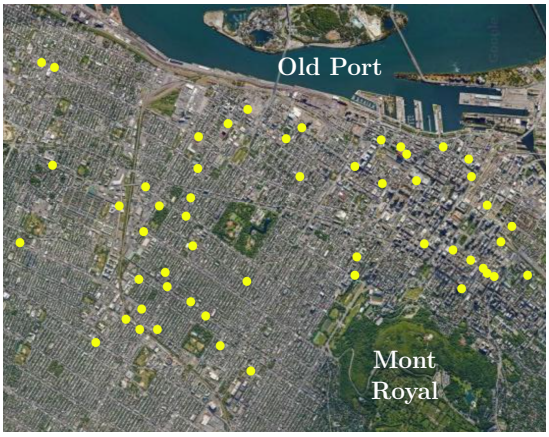

Figure 5.9 Cluster of BIXI stations located in Montreal (Canada)

We follow the approach of [70] to cluster the stations according to their trip behaviour using k-means. We then select a cluster around the downtown and plateau areas in which the total number of rentals is approximately the same as the total number of returns. This cluster has 53 stations, including several city center stations and therefore contains work-related trips. Given that the distances between stations inside the cluster are limited, vehicles have sufficient time available to relocate and rebalance bikes within each time-period. The stations in the selected cluster are visualized in Figure 5.9.

Results on Real-world data

Given that, on synthetic data, DROB-I within rolling planning outperformed DROB-T on lost rental under all the predictions and consistently had swift computing times, we here focus on comparing the performance of DROB-I and DROB-LD. Experiments are carried out in a rolling planning, which aligns with practice and accommodates the need for swift runtime, while also allowing for real-time system updates.

The results are visualized in Figure 5.10. Detailed results for each day can be found in Appendix D. DROB-I performs better on both lost rental and total lost demand, while the lost return is higher than for DROB-LD. Note that lower lost rental demand (and therefore more succesful trips) also results in more return demand, which explains that DROB-I suffers from a slightly higher lost return. Such observations align with previous results on synthetic data. The results on real-world data confirm the benefits of model DROB-I, providing robustness to the station inventories and reducing the total lost demand. In terms of computing time, both models can be solved to optimality within 1 minute.

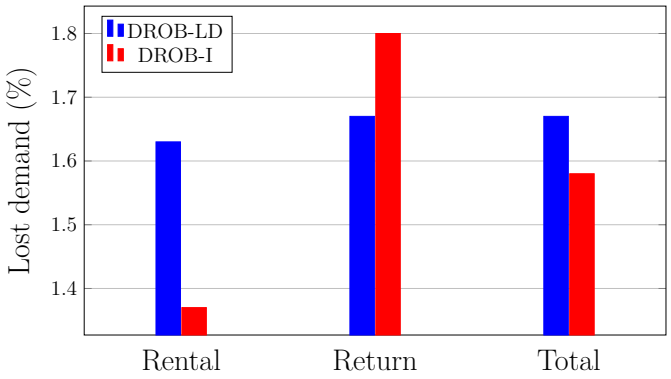

Figure 5.10 Lost demand of DROB-I and DROB-LD on a cluster from BIXI

In contrast to the results on synthetic data, the improvement provided by DROB-I is not impressive. Such smaller improvement may be explained by the fact that the here-considered real-world data only contains successful trips. We also carried out static planning for the DROB-LD and DROB-I. However, under static planning, DROB-LD runs out of memory for many cases, given that the model contains more time-periods. The total lost demand of DROB-I under static planning has been 1.98%, shortly higher than the one under rolling planning (1.58%), highlighting once again the benefits of the rolling planning for DROB-I.

5.5 Conclusions

In this work, we have proposed two objective functions for multi-period rebalancing models for Bike-sharing systems, DROB-I and DROB-T, incorporating inventory intervals and target inventories. The resulting models provide an alternative to classical models minimizing unmet demand and are particularly suitable for BSSs operators that use (often manually computed) inventory intervals and targets to guide their rebalancing process.

Our work evaluates the entire pipeline required for an automatized and data-driven rebalancing process. Instead of relying on manual input, we estimate rental and return demand for each hour and station in a data-driven fashion, using a machine-learning model that has been shown to provide reasonably accurate results based on historical data related to time, weather and user trips. Inventory intervals and targets are then derived such that they maximize the desired service-level.

Our empirical analysis explore the capability of the proposed planning solutions to meet customer demands under three key characteristics of the planning process: the used optimization model, the employed reoptimization mode, and the impact of highly accurate (or inaccurate) demand predictions. The obtained planning solutions are then evaluated within

a fine-grained minute-by-minute discrete-event simulator. A series of experiments on synthetic data allows for three key conclusions: First, our proposed models exhibit remarkable robustness compared to DROB-LD, the classical model minimizing unmet demand. DROB-T leads to a reduction in lost demand of up to 34%, while DROB-I decreases lost demand by up to 28%. Second, such robustness is also observed when demand predictions are less accurate, as these models introduce a conservative buffer into the station inventories, capable of better dealing with stochastic demand fluctuations. Third, there is a pronounced benefit in reoptimizing the rebalancing decisions throughout the planning as opposed to executing the optimization model only once and implementing a static planning solution for the entire planning horizon. Allowing for updated system information, the improvement via a rolling or folding planning has been found to be consistently in the order of 15-20% as opposed to static planning for DROB-LD. For DROB-I and DROB-T, the improvement tends to be higher than 30%, clearly indicating the benefits of such additional reoptimization effort. Finally, the benefits of our proposed models observed on synthetic data are also verified on a case study with real-world data from a BIXI Montreal.

Our approach may be highly attractive to system operators, not only due to their superior performance, but also due to their fit within the existing decision-making processes, as inventory intervals and targets are often used concepts in practice. In addition, we hope that the here proposed weather generator inspires future research to evaluate planning approaches in a more complex and realistic manner. Given the benefits of the here proposed models, we believe that the development of tailored solution methods (such as mathematical decomposition methods) may be promising research directions, highly useful for both academia and practitioners to approach rebalancing in large-scale station networks.

CHAPTER 6 GENERAL DISCUSSION

This thesis contributes to data-driven rebalancing strategies for BSSs, focusing on the deployment of rebalancing recommendations and planning based on inventory intervals and target inventory values. A summary of the proposed models is detailed in Section 6.1, followed by their applicability and limitations in Section 6.2.

6.1 Summary of Works

Studies that tackle the rebalancing problem at a tactical level tend to focus solely on the computation of inventory intervals and target inventory values. The model introduced in Chapter 3 presents a novel approach in which inventory intervals and target inventory values are combined with data-driven prioritization strategies to automatically select a subset of unbalanced stations to be rebalanced at every hour. Reducing the set of stations to be rebalanced is important, given that BSSs have limited resources to redistribute bikes in the network, making it impractical to rebalance all unbalanced stations in the system at every hour. Moreover, the integration of the proposed model into real-world BSSs is facilitated by its minimal computational needs – an advantage compared to rebalancing models on the operational level. Also, the proposed prioritization strategy and the configuration of the inventory intervals can be modified to adjust the rebalancing recommendations depending on the preferences of the operator. For example, if the operator aims to maintain high service levels, they can opt for narrow inventory intervals, ensuring that the stations trigger a rebalancing alert whenever their inventories are slightly unbalanced. Similarly, by changing the prioritization strategy, the operator can control which stations are prioritized during the rebalancing process. The obtained results show that the proposed prioritization strategies outperform the current prioritization practiced by BIXI, managing to reduce both lost demand and rebalancing operations simultaneously in the majority of the experiments. This demonstrates that selecting the stations to be rebalanced according to the predicted demand can drastically improve the performance of the rebalancing operations while respecting the BSS rebalancing capacity.

Also on the tactical level, the model introduced in Chapter 4 tackles the rebalancing problem of bimodal shared bikes, i.e., regular and e-bikes. To the best of our knowledge, literature has not yet explored this specific problem. Existing studies, considering a single type of bike in the system, cannot be easily adapted to bimodal BSSs, as they may generate unrealistic rebalancing recommendations – e.g., the target inventory value of both demands can exceed

the total capacity of the station. Consequently, our model stands out for offering tailored rebalancing recommendations for regular and electric bikes in a dock-based bimodal BSS, considering the maximum capacity of the station as well as where both bike types share the available docks. Furthermore, the flexibility to adjust the inventory intervals for regular and electric bikes separately represents an important advantage given that each demand type can be present quite different patterns. Thus, this study emphasizes the importance of understanding user preferences and adapting the rebalancing strategy accordingly. In addition, the results indicate that encouraging users to consider alternative bikes, if their first choice of bike is unavailable, can enhance service levels. Our proposed model reduces the total lost demand in all simulated scenarios, especially for electric bikes, demonstrating better adaptability to changing demands and commuters' preferences.

In Chapter 5, we delve into optimization models for BSS that actually obtain rebalancing routes aimed at minimizing the deviation from inventory intervals and target inventory values. While most studies presented in the literature assume predefined inventory intervals and target inventory values, we employ a predictive model that considers weather, temporal, and historical trip data to predict trip demand. Based on these demand predictions, we compute the inventory intervals and target inventory values. We then compare our models with a well-known strategy that optimizes rebalancing routes minimizing lost demand. This comparison is particularly significant as most proposed optimization models typically focus on evaluating the effectiveness of a single objective function. Our study, however, compares the performance of three distinct objective functions under three distinct reoptimization modes (static planning, rolling planning and folding planning). The obtained results indicate that the performance of the studied BSS tends to be more stable when the rebalancing routes are optimized based on the inventory intervals and target inventory values, whereas the performance of the baseline model is more sensitive to demand fluctuations. Thus, our proposed models present a more robust strategy to mitigate lost demand by ensuring a buffer to the stations' inventory. Moreover, the results show that updating system status (in our case, the inventory of the stations and the demand prediction) when reoptimizing the rebalancing strategies always leads to improvement of the BSS performance, reducing lost demand by up to 36% when compared to the same rebalancing strategies that do not update the system status.

6.2 Applicability and Limitations

The models presented in Chapters 3 and 4 have relatively low computational requirements and require solely a database to fuel the predictive algorithm, making them easy to imple-

ment in real-world BSSs. This is a significant advantage, as some mathematically complex models require significant computing power and time, which can be a constraint for certain operators. Therefore, while other approaches may be theoretically optimal, they may not be operationally feasible in practice. Furthermore, our proposed models allow scaling to BSS of various sizes. However, it is important to note that the models presented in Chapters 3 and 4 primarily focus on providing rebalancing recommendations on a tactical level, without incorporating the routing aspect. As a result, they may require operator intervention to design the rebalancing routes.

An important limitation of the data-driven model presented in Chapter 4 for shared BSSs with regular and electric bikes is that it disregards the charging level of batteries in e-bikes over the simulated period, considering that e-bikes are always available for renting. In reality, uncharged e-bikes are not available to be rented as the pedal-assist cannot work properly. This indeed affects our simulation as it reduces the amount of rebalancing alerts triggered by stations with uncharged e-bikes. Consequently, less lost demand is very likely computed within the simulation in comparison to real-world scenarios.

Finally, the optimization models proposed in Chapter 5 obtain rebalancing routes that consider the capacity of the stations and trucks, trip duration between stations, and location of the stations. Therefore, these optimization models drastically reduce the need for manual interventions. Nonetheless, under unforeseen circumstances, such as traffic jams, accidents, and blocked roads, external interference is still required to analyze the feasibility of such rebalancing routes. It is also worth noting that the complexity of DROB-I and DROB-T, with their numerous variables and constraints, requires substantial computational effort, which limits their applicability to small or medium-scale BSS or a cluster of stations.

CHAPTER 7 CONCLUSION

Planning BSS rebalancing is a complex task, given the many factors that have to be taken into account, such as station demand for the upcoming hours, the number of docks at each station, the number of bikes available in the system and the number of trucks and their capacity. In this context, data-driven techniques can provide a sophisticated approach to managing these complexities. By analyzing historical trip demand and its correlation with weather and temporal features, data-driven models can forecast demand at a station level and, subsequently, estimate a rebalancing schedule that maximizes the successful trips in the system. This strategy allows rebalancing recommendations to be made in order to prevent stations from reaching empty or full status – which would likely result in lost demand later. Furthermore, given the limited number of trucks available for rebalancing, a data-driven model can systematically indicate which stations have the highest need for intervention, optimizing the BSS's resources.

In Chapters 3 and 4, we compared our proposed data-driven rebalancing strategies with the current rebalancing strategy used by BIXI operators. The results consistently demonstrated an enhancement in service level when employing data-driven rebalancing strategies, improving the availability of bikes and docks in the system. Furthermore, our analysis revealed that data-driven rebalancing strategies are also able to reduce the total number of rebalancing operations in most of the performed tests. These outcomes highlight that this approach can prevent unnecessary operations, avoiding stations that would naturally return to a balanced state, and focus the rebalancing resources on the most critical stations.

The rebalancing recommendations proposed in this thesis are based on two mechanisms, inventory intervals and target inventory values. These mechanisms have demonstrated practical utility, as evidenced by their implementation in some real-world BSSs, such as BIXI. In addition, the use of inventory intervals makes it easier to adjust the rebalancing recommendations according to the strategy chosen by the operator, leading to rebalancing operations that aim to achieve the desired service level for the network. In Chapter 5, we demonstrate the effectiveness of these mechanisms when used within rebalancing optimization models that provide optimized routes along with rebalancing inventory management actions (i.e., add/drop of bikes at the stations). The resulting models are compared with a traditional rebalancing model found in the literature that exclusively minimizes lost demand. Our results indicate that the proposed models offer more robustness, reducing lost demand even in scenarios where demand prediction is less precise. This outcome underscores the advantage of

our approach in maintaining reasonable bike levels at the stations regardless of the accuracy of demand prediction.

7.1 Future Work

The rebalancing models presented in this thesis benefit from a regression model to estimate future demand and propose rebalancing recommendations accordingly. Given that bike demand is heavily influenced by seasonal changes, weather conditions, and urban events, the regression models must continuously update the training dataset to learn recent demand trends and ensure an accurate prediction. A possible area to be explored is determining the optimal frequency for updating the training dataset of the regression model. This aspect is critical to maintain the desired level of accuracy in demand prediction and ensure the effectiveness of the rebalancing recommendations proposed by our models.

The rebalancing models presented in this thesis have primarily adopted an operator-based approach, where the logistics of rebalancing are directly planned and executed by the operators. Nevertheless, many large-scale BSS employ a combination of operator-based and user-based approaches (where users receive incentives to allocate bikes to high rental demand stations) to improve their service level. Therefore, a logical direction for future research is to develop a comprehensive model that considers both operator-based and user-based approaches within the system. Such a model could provide a more holistic and integrated solution, leveraging the active involvement of both operators and users to optimize the rebalancing process, improve system efficiency, and enhance user satisfaction.

The prioritization strategies introduced in Chapter 3 offer a promising avenue for enhancing the practicality of optimization models. By generating rebalancing routes that focus on a subset of high-priority stations, rather than the entire network, these prioritization strategies effectively reduce the size of the problem and, hence, contribute to the traceability of the optimization models. For this reason, a possible next step is to integrate the prioritization strategies into the proposed models DROB-I and DROB-T. This integration could be particularly effective within a rolling planning, in which the models would dynamically design rebalancing routes at each stage of reoptimization, each time considering a different subset of prioritized stations. Such an approach would bring a novel perspective to rebalancing operations, focusing resources where they are most needed at any given time. To further advance this research, it would be valuable to conduct a comparative analysis of this methodology against the traditional Cluster-First Route-Second strategy, which is commonly used to enhance the feasibility of rebalancing optimization models. A comprehensive comparison of these two approaches would provide deep insights into their respective strengths

and limitations.

REFERENCES

[1] H. Akova, S. Hulagu, and H. B. Celikoglu, "Static bike repositioning problem with heterogeneous distribution characteristics in bike sharing systems," *Transportation Research Procedia*, vol. 62, pp. 205–212, 2022.

[2] R. Alvarez-Valdes, J. M. Belenguer, E. Benavent, J. D. Bermudez, F. Muñoz, E. Vercher, and F. Verdejo, "Optimizing the level of service quality of a bike-sharing system," *Omega*, vol. 62, pp. 163–175, 2016.

[3] E. Basak, R. Al Balawi, S. Fatemi, and A. Tafti, "When crisis hits: Bike-sharing platforms amid the covid-19 pandemic," *Plos one*, vol. 18, no. 4, p. e0283603, 2023.

[4] M. Benchimol, P. Benchimol, B. Chappert, A. De La Taille, F. Laroche, F. Meunier, and L. Robinet, "Balancing the stations of a self service "bike hire" system," *RAIRO-Operations Research-Recherche Opérationnelle*, vol. 45, no. 1, pp. 37–61, 2011.

[5] P. Borgnat, P. Abry, P. Flandrin, C. Robardet, J.-B. Rouquier, and E. Fleury, "Shared bicycles in a city: A signal processing and data analysis perspective," *Advances in Complex Systems*, vol. 14, no. 03, pp. 415–438, 2011.

[6] N. Boufidis, A. Nikiforiadis, K. Chrysostomou, and G. Aifadopoulou, "Development of a station-level demand prediction and visualization tool to support bike-sharing systems' operators," *Transportation Research Procedia*, vol. 47, pp. 51–58, 2020.

[7] J. Brinkmann, M. W. Ulmer, and D. C. Mattfeld, "Short-term strategies for stochastic inventory routing in bike sharing systems," *Transportation Research Procedia*, vol. 10, pp. 364–373, 2015.

[8] ——, "Inventory routing for bike sharing systems," *Transportation research procedia*, vol. 19, pp. 316–327, 2016.

[9] ——, "The multi-vehicle stochastic-dynamic inventory routing problem for bike sharing systems," *Business Research*, vol. 13, no. 1, pp. 69–92, 2020.

[10] T. Bulhões, A. Subramanian, G. Erdoğan, and G. Laporte, "The static bike relocation problem with multiple vehicles and visits," *European Journal of Operational Research*, vol. 264, no. 2, pp. 508–523, 2018.

[11] G. C. Calafiore, C. Bongiorno, and A. Rizzo, "A robust mpc approach for the rebalancing of mobility on demand systems," *Control Engineering Practice*, vol. 90, pp. 169–181, 2019.

[12] K. B. Campbell and C. Brakewood, "Sharing riders: How bikesharing impacts bus ridership in new york city," *Transportation Research Part A: Policy and Practice*, vol. 100, pp. 264–282, 2017.

[13] D. Chemla, F. Meunier, and R. W. Calvo, "Bike sharing systems: Solving the static rebalancing problem," *Discrete Optimization*, vol. 10, no. 2, pp. 120–146, 2013.

[14] D. Chemla, F. Meunier, T. Pradeau, R. W. Calvo, and H. Yahiaoui, "Self-service bike sharing systems: simulation, repositioning, pricing," Centre d'Enseignement et de Recherche en Mathématiques et Calcul Scientifique - CERMICS, Tech. Rep. hal-00824078f, 2013.

[15] L. Chen, D. Zhang, L. Wang, D. Yang, X. Ma, S. Li, Z. Wu, G. Pan, T.-M.-T. Nguyen, and J. Jakubowicz, "Dynamic cluster-based over-demand prediction in bike sharing systems," in *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, 2016, pp. 841–852.

[16] P.-C. Chen, H.-Y. Hsieh, K.-W. Su, X. K. Sigalingging, Y.-R. Chen, and J.-S. Leu, "Predicting station level demand in a bike-sharing system using recurrent neural networks," *IET Intelligent Transport Systems*, vol. 14, no. 6, pp. 554–561, 2020.

[17] Q. Chen, C. Fu, N. Zhu, S. Ma, and Q.-C. He, "A target-based optimization model for bike-sharing systems: From the perspective of service efficiency and equity," *Transportation Research Part B: Methodological*, vol. 167, pp. 235–260, 2023. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S019126152200203X

[18] Z. Chen, Y. Hu, J. Li, and X. Wu, "Optimal deployment of electric bicycle sharing stations: model formulation and solution technique," *Networks and Spatial Economics*, vol. 20, no. 1, pp. 99–136, 2020.

[19] Y. Chumin, O. O'Brien, P. DeMaio, R. Rabello, S. Chou, and T. Benicchio, "The meddin bike-sharing world map report," https://bikesharingworldmap.com/reports/bswm_mid2021report.pdf, 2021, accessed: 2023-08-04.

[20] C. Contardo, C. Morency, and L.-M. Rousseau, *Balancing a dynamic public bike-sharing system*. Cirrelt Montreal, Canada, 2012, vol. 4.

[21] S. Datner, T. Raviv, M. Tzur, and D. Chemla, "Setting inventory levels in a bike sharing network," *Transportation Science*, vol. 53, no. 1, pp. 62–76, 2019.

[22] M. Dell'Amico, E. Hadjicostantinou, M. Iori, and S. Novellani, "The bike sharing rebalancing problem: Mathematical formulations and benchmark instances," *Omega*, vol. 45, pp. 7–19, 2014.

[23] P. DeMaio, "Bike-sharing: History, impacts, models of provision, and future," *Journal of public transportation*, vol. 12, no. 4, p. 3, 2009.

[24] L. Di Gaspero, A. Rendl, and T. Urli, "Constraint-based approaches for balancing bike sharing systems," in *International Conference on Principles and Practice of Constraint Programming*. Springer, 2013, pp. 758–773.

[25] W. El-Assi, M. Salah Mahmoud, and K. Nurul Habib, "Effects of built environment and weather on bike sharing demand: a station level analysis of commercial bike sharing in toronto," *Transportation*, vol. 44, pp. 589–613, 2017.

[26] G. Erdoğan, M. Battarra, and R. W. Calvo, "An exact algorithm for the static rebalancing problem arising in bicycle sharing systems," *European Journal of Operational Research*, vol. 245, no. 3, pp. 667–679, 2015.

[27] G. Erdoğan, G. Laporte, and R. Wolfler Calvo, "The static bicycle relocation problem with demand intervals," *European Journal of Operational Research*, vol. 238, no. 2, pp. 451–457, 2014.

[28] E. Eren and V. E. Uz, "A review on bike-sharing: The factors affecting bike-sharing demand," *Sustainable cities and society*, vol. 54, p. 101882, 2020.

[29] H. M. Espegren, J. Kristianslund, H. Andersson, and K. Fagerholt, "The static bicycle repositioning problem-literature survey and new formulation," in *Computational Logistics: 7th International Conference, ICCL 2016, Lisbon, Portugal, September 7-9, 2016, Proceedings 7*. Springer, 2016, pp. 337–351.

[30] A. Faghih-Imani and N. Eluru, "Incorporating the impact of spatio-temporal interactions on bicycle sharing system demand: A case study of new york citibike system," *Journal of Transport Geography*, vol. 54, pp. 218–227, 2016.

[31] S. Feng, H. Chen, C. Du, J. Li, and N. Jing, "A hierarchical demand prediction method with station clustering for bike sharing system," in *2018 IEEE Third International Conference on Data Science in Cyberspace (DSC)*. IEEE, 2018, pp. 829–836.

[32] I. A. Forma, T. Raviv, and M. Tzur, "A 3-step math heuristic for the static repositioning problem in bike-sharing systems," *Transportation research part B: methodological*, vol. 71, pp. 230–247, 2015.

[33] C. Fricker and N. Gast, "Incentives and redistribution in homogeneous bike-sharing systems with stations of finite capacity," *Euro journal on transportation and logistics*, vol. 5, no. 3, pp. 261–291, 2016.

[34] J. Froehlich, J. Neumann, N. Oliver *et al.*, "Sensing and predicting the pulse of the city through shared bicycling." in *IJCAI*, vol. 9, no. Jul, 2009, pp. 1420–1426.

[35] T. Fukushige, D. T. Fitch, and S. Handy, "Can an incentive-based approach to rebalancing a dock-less bike-share system work? evidence from sacramento, california," *Transportation Research Part A: Policy and Practice*, vol. 163, pp. 181–194, 2022.

[36] C. Gallop, C. Tse, and J. Zhao, "A seasonal autoregressive model of vancouver bicycle traffic using weather variables," *i-Manager's Journal on Civil Engineering*, vol. 1, no. 4, p. 9, 2011.

[37] D. Gammelli, Y. Wang, D. Prak, F. Rodrigues, S. Minner, and F. C. Pereira, "Predictive and prescriptive performance of bike-sharing demand forecasts for inventory management," *Transportation Research Part C: Emerging Technologies*, vol. 138, p. 103571, 2022.

[38] ——, "Predictive and prescriptive performance of bike-sharing demand forecasts for inventory management," *Transportation Research Part C: Emerging Technologies*, vol. 138, p. 103571, 2022.

[39] K. Gebhart and R. B. Noland, "The impact of weather conditions on bikeshare trips in washington, dc," *Transportation*, vol. 41, no. 6, pp. 1205–1225, 2014.

[40] D. K. George and C. H. Xia, "Fleet-sizing and service availability for a vehicle rental system via closed queueing networks," *European Journal of Operational Research*, vol. 211, no. 1, pp. 198–207, 2011.

[41] S. Ghosh, P. Varakantham, Y. Adulyasak, and P. Jaillet, "Dynamic redeployment to counter congestion or starvation in vehicle sharing systems," in *Twenty-Fifth International Conference on Automated Planning and Scheduling*, 2015.

[42] S. Ghosh, M. Trick, and P. Varakantham, "Robust repositioning to counter unpredictable demand in bike sharing systems," 2016.

[43] S. Ghosh, P. Varakantham, Y. Adulyasak, and P. Jaillet, "Dynamic repositioning to reduce lost demand in bike sharing systems," *Journal of Artificial Intelligence Research*, vol. 58, pp. 387–430, 2017.

[44] S. Ghosh, J. Y. Koh, and P. Jaillet, "Improving customer satisfaction in bike sharing systems through dynamic repositioning," in *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, 2019, pp. 5864–5870.

[45] M. D. Gleditsch, K. Hagen, H. Andersson, S. J. Bakker, and K. Fagerholt, "A column generation heuristic for the dynamic bicycle rebalancing problem," *European Journal of Operational Research*, 2022.

[46] R. C. Hampshire and L. Marla, "An analysis of bike sharing usage: Explaining trip generation and attraction from observed demand," in *91st Annual meeting of the transportation research board, Washington, DC*, 2012, pp. 12–2099.

[47] S. Hossain, M. A. Islam, and M. S. Akther, "Covid-19 impact on travel and work habits of office workers in bangladesh," *Transportation Engineering*, vol. 11, p. 100162, 2023.

[48] K. Hosseini, A. Stefaniec, M. O'Mahony, and B. Caulfield, "Optimising shared electric mobility hubs: Insights from performance analysis and factors influencing riding demand," *Case Studies on Transport Policy*, vol. 13, p. 101052, 2023.

[49] R. Hu, Z. Zhang, X. Ma, and Y. Jin, "Dynamic rebalancing optimization for bike-sharing system using priority-based moea/d algorithm," *IEEE Access*, vol. 9, pp. 27 067–27 084, 2021.

[50] J. Huang, H. Sun, H. Li, L. Huang, A. Li, and X. Wang. (2020) Central station-based demand prediction for determining target inventory in a bike-sharing system.

[51] J. Huang, Q. Tan, H. Li, A. Li, and L. Huang, "Monte carlo tree search for dynamic bike repositioning in bike-sharing systems," *Applied Intelligence*, pp. 1–16, 2022.

[52] P. Hulot, "Towards station-level demand prediction for effective rebalancing in bike-sharing systems," Master's thesis, École Polytechnique de Montréal, 2018.

[53] P. Hulot, D. Aloise, and S. D. Jena, "Towards station-level demand prediction for effective rebalancing in bike-sharing systems," in *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2018, pp. 378–386.

[54] G. Héctor R., L. Rafael, and A. Ramirez-Nafarrate, "A simulation-optimization study of the inventory of a bike-sharing system: The case of mexico city ecobici's system," *Case Studies on Transport Policy*, vol. 9, no. 3, pp. 1059–1072, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2213624X21000791

[55] S. Ji, C. R. Cherry, L. D. Han, and D. A. Jordan, "Electric bike sharing: simulation of user demand and system availability," *Journal of Cleaner Production*, vol. 85, pp. 250–257, 2014.

[56] W. Jia, Y. Tan, L. Liu, J. Li, H. Zhang, and K. Zhao, "Hierarchical prediction based on two-level gaussian mixture model clustering for bike-sharing system," *Knowledge-Based Systems*, vol. 178, pp. 84–97, 2019.

[57] N. Jian and S. G. Henderson, "An introduction to simulation optimization," in *2015 Winter Simulation Conference (WSC)*, 2015, pp. 1780–1794.

[58] Y. Jin, C. Ruiz, and H. Liao, "A simulation framework for optimizing bike rebalancing and maintenance in large-scale bike-sharing systems," *Simulation Modelling Practice and Theory*, vol. 115, p. 102422, 2022.

[59] R. Julio and A. Monzon, "Long term assessment of a successful e-bike-sharing system. key drivers and impact on travel behaviour," *Case Studies on Transport Policy*, 2022.

[60] A. Kabra, E. Belavina, and K. Girotra, "Bike-share systems: Accessibility and availability," *Management Science*, vol. 66, no. 9, pp. 3803–3824, 2020.

[61] K. Kim, "Investigation on the effects of weather and calendar events on bike-sharing according to the trip patterns of bike rentals of stations," *Journal of transport geography*, vol. 66, pp. 309–320, 2018.

[62] C. Kloimüllner, P. Papazek, B. Hu, and G. R. Raidl, "Balancing bicycle sharing systems: an approach for the dynamic case," in *European Conference on Evolutionary Computation in Combinatorial Optimization*. Springer, 2014, pp. 73–84.

[63] ——, "A cluster-first route-second approach for balancing bicycle sharing systems," in *Computer Aided Systems Theory–EUROCAST 2015: 15th International Conference, Las Palmas de Gran Canaria, Spain, February 8-13, 2015, Revised Selected Papers 15*. Springer, 2015, pp. 439–446.

[64] B. Legros. (2019) Dynamic repositioning strategy in a bike-sharing system; how to prioritize and how to rebalance a bike station.

[65] G. Li, N. Cao, P. Zhu, Y. Zhang, Y. Zhang, L. Li, Q. Li, and Y. Zhang, "Towards smart transportation system: A case study on the rebalancing problem of bike sharing system based on reinforcement learning," *Journal of Organizational and End User Computing (JOEUC)*, vol. 33, no. 3, pp. 35–49, 2021.

[66] X. Li, Y. Xu, X. Zhang, W. Shi, Y. Yue, and Q. Li, "Improving short-term bike sharing demand forecast through an irregular convolutional neural network," *Transportation research part C: emerging technologies*, vol. 147, p. 103984, 2023.

[67] Y. Li and Y. Zheng, "Citywide bike usage prediction in a bike-sharing system," *IEEE Transactions on Knowledge and Data Engineering*, vol. 32, no. 6, pp. 1079–1091, 2019.

[68] Y. Li, Y. Zheng, H. Zhang, and L. Chen, "Traffic prediction in a bike-sharing system," in *Proceedings of the 23rd SIGSPATIAL international conference on advances in geographic information systems*, 2015, pp. 1–10.

[69] Y. Li, Y. Zheng, and Q. Yang, "Dynamic bike reposition: A spatio-temporal reinforcement learning approach," in *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2018, pp. 1724–1733.

[70] J. Liang, S. D. Jena, and A. Lodi, "Dynamic rebalancing optimization for bike-sharing systems: A modeling framework and empirical comparison," GERAD, HEC Montréal, Canada, Les Cahiers du GERAD G–2023–47, 2023.

[71] S. Lin, F. Y. Chen, Y. Li, and Z.-J. M. Shen, "Dynamic inventory control with covariates: Risk constraints, regularization, and folding-horizon plan," *Yanzhi and Shen, Zuo-Jun Max, Dynamic Inventory Control with Covariates: Risk Constraints, Regularization, and Folding-horizon Plan (February 14, 2022)*, 2022.

[72] J. Liu, L. Sun, W. Chen, and H. Xiong, "Rebalancing bike sharing systems: A multi-source data smart optimization," in *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, 2016, pp. 1005–1014.

[73] M. Lowalekar, P. Varakantham, S. Ghosh, S. D. Jena, and P. Jaillet, "Online repositioning in bike sharing systems," in *Twenty-Seventh International Conference on Automated Planning and Scheduling*, 2017.

[74] Á. Lozano, J. F. De Paz, G. Villarrubia Gonzalez, D. H. D. L. Iglesia, and J. Bajo, "Multi-agent system for demand prediction and trip visualization in bike sharing systems," *Applied Sciences*, vol. 8, no. 1, p. 67, 2018.

[75] C.-C. Lu, "Robust multi-period fleet allocation models for bike-sharing systems," *Networks and Spatial Economics*, vol. 16, no. 1, pp. 61–82, 2016.

[76] E. H.-C. Lu and Z.-Q. Lin, "Rental prediction in bicycle-sharing system using recurrent neural network," *IEEE Access*, vol. 8, pp. 92 262–92 274, 2020.

[77] Y. Lu, U. Benlic, and Q. Wu, "An effective memetic algorithm for the generalized bike-sharing rebalancing problem," *Engineering Applications of Artificial Intelligence*, vol. 95, p. 103890, 2020.

[78] J. F. Magee, *Decision trees for decision making*. Harvard Business Review Brighton, MA, USA, 1964.

[79] L. M. Martinez, L. Caetano, T. Eiró, and F. Cruz, "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*, vol. 54, pp. 513–524, 2012.

[80] M. C. Martins Silva, D. Aloise, and S. D. Jena, "Data-driven prioritization strategies for inventory rebalancing in bike-sharing systems," Groupe d'études et de recherche en analyse des décisions, GERAD, Montréal QC H3T 2A7, Canada, Les Cahiers du GERAD G-2023-30, Aug. 2023. [Online]. Available: https://www.gerad.ca/en/papers/G-2023-30

[81] ——, "Towards effective rebalancing of bike-sharing systems with regular and electric bikes," Groupe d'études et de recherche en analyse des décisions, GERAD, Montréal QC H3T 2A7, Canada, Les Cahiers du GERAD G-2023-31, Aug. 2023. [Online]. Available: https://www.gerad.ca/en/papers/G-2023-31

[82] K. Mellou and P. Jaillet, "Dynamic resource redistribution and demand estimation: An application to bike sharing systems," *SSRN Electron*, pp. 1–58, 2019.

[83] P. Mrazovic, J. L. Larriba-Pey, and M. Matskin, "A deep learning approach for estimating inventory rebalancing demand in bicycle sharing systems," in *2018 IEEE 42nd Annual Computer Software and Applications Conference (COMPSAC)*, vol. 2. IEEE, 2018, pp. 110–115.

[84] R. Nair and E. Miller-Hooks, "Fleet management for vehicle sharing operations," *Transportation Science*, vol. 45, no. 4, pp. 524–540, 2011.

[85] Z. Niu and L. Chai, "Carbon emission reduction by bicycle-sharing in china," *Energies*, vol. 15, no. 14, p. 5136, 2022.

[86] P. Nunes, A. Moura, and J. Santos, "Solving the multi-objective bike routing problem by meta-heuristic algorithms," *International Transactions in Operational Research*, 2022.

[87] O. O'Brien, P. DeMaio, R. Rabello, S. Chou, and T. Benicchio, "The meddin bike-sharing world map report," https://bikesharingworldmap.com/reports/bswm_mid2022report.pdf, 2022, accessed: 2023-08-04.

[88] A. Pal and Y. Zhang, "Free-floating bike sharing: Solving real-life large-scale static rebalancing problems," *Transportation Research Part C: Emerging Technologies*, vol. 80, pp. 92–116, 2017.

[89] Y. Pan, R. C. Zheng, J. Zhang, and X. Yao, "Predicting bike sharing demand using recurrent neural networks," *Procedia computer science*, vol. 147, pp. 562–566, 2019.

[90] P. Papazek, G. R. Raidl, M. Rainer-Harbach, and B. Hu, "A pilot/vnd/grasp hybrid for the static balancing of public bicycle sharing systems," in *International Conference on Computer Aided Systems Theory*. Springer, 2013, pp. 372–379.

[91] P. Papazek, C. Kloimüllner, B. Hu, and G. R. Raidl, "Balancing bicycle sharing systems: an analysis of path relinking and recombination within a grasp hybrid," in *International Conference on Parallel Problem Solving from Nature*. Springer, 2014, pp. 792–801.

[92] F. Pase, F. Chiariotti, A. Zanella, and M. Zorzi, "Bike sharing and urban mobility in a post-pandemic world," *IEEE Access*, vol. 8, pp. 187 291–187 306, 2020.

[93] E. Possani and E. Castillo, "Optimizing the inventory and routing decisions in a bike-sharing system: A linear programming and stochastic approach," *Case Studies on Transport Policy*, vol. 9, no. 4, pp. 1495–1502, 2021.

[94] J. Pucher, R. Buehler, D. R. Bassett, and A. L. Dannenberg, "Walking and cycling to health: a comparative analysis of city, state, and international data," *American journal of public health*, vol. 100, no. 10, pp. 1986–1992, 2010.

[95] M. Rainer-Harbach, P. Papazek, G. R. Raidl, B. Hu, and C. Kloimüllner, "Pilot, grasp, and vns approaches for the static balancing of bicycle sharing systems," *Journal of Global Optimization*, vol. 63, no. 3, pp. 597–629, 2015.

[96] A. A. Ramesh, S. P. Nagisetti, N. Sridhar, K. Avery, and D. Bein, "Station-level demand prediction for bike-sharing system," in *2021 IEEE 11th Annual Computing and Communication Workshop and Conference (CCWC)*. IEEE, 2021, pp. 0916–0921.

[97] T. Raviv and O. Kolka, "Optimal inventory management of a bike-sharing station," *IIE Transactions*, vol. 45, no. 10, pp. 1077–1093, 2013.

[98] T. Raviv, M. Tzur, and I. A. Forma, "Static repositioning in a bike-sharing system: models and solution approaches," *EURO Journal on Transportation and Logistics*, vol. 2, no. 3, pp. 187–229, 2013.

[99] Y. Ren, L. Meng, F. Zhao, C. Zhang, H. Guo, Y. Tian, W. Tong, and J. W. Sutherland, "An improved general variable neighborhood search for a static bike-sharing rebalancing problem considering the depot inventory," *Expert Systems with Applications*, vol. 160, p. 113752, 2020.

[100] R. A. Rixey, "Station-level forecasting of bikesharing ridership: Station network effects in three us systems," *Transportation research record*, vol. 2387, no. 1, pp. 46–55, 2013.

[101] K. Roshan Zamir, "Dynamic repositioning for bikesharing systems," Ph.D. dissertation, 2020.

[102] C. Rudloff and B. Lackner, "Modeling demand for bicycle sharing system–neighboring stations as a source for demand and a reason for structural breaks," in *Transportation Research Board Annual Meeting*, 2013.

[103] S. Ruffieux, N. Spycher, E. Mugellini, and O. Abou Khaled, "Real-time usage forecasting for bike-sharing systems: A study on random forest and convolutional neural network applicability," in *2017 intelligent systems conference (intellisys)*. IEEE, 2017, pp. 622–631.

[104] V. Sathishkumar, P. Jangwoo, and C. Yongyun, "Using data mining techniques for bike sharing demand prediction in metropolitan city," *Computer Communications*, vol. 153, pp. 353–366, 2020.

[105] H. Sayarshad, S. Tavassoli, and F. Zhao, "A multi-periodic optimization formulation for bike planning and bike utilization," *Applied Mathematical Modelling*, vol. 36, no. 10, pp. 4944–4951, 2012.

[106] J. Schuijbroek, R. C. Hampshire, and W.-J. Van Hoeve, "Inventory rebalancing and vehicle routing in bike sharing systems," *European Journal of Operational Research*, vol. 257, no. 3, pp. 992–1004, 2017.

[107] Y. Seo, "A dynamic rebalancing strategy in public bicycle sharing systems based on real-time dynamic programming and reinforcement learning," Ph.D. dissertation, Doctoral dissertation. Seoul National University, South Korea, 2020.

[108] M. E. Shaik and S. Ahmed, "An overview of the impact of covid-19 on road traffic safety and travel behavior," *Transportation Engineering*, vol. 9, p. 100119, 2022.

[109] J. Shu, M. C. Chou, Q. Liu, C.-P. Teo, and I.-L. Wang, "Models for effective deployment and redistribution of bicycles within public bicycle-sharing systems," *Operations Research*, vol. 61, no. 6, pp. 1346–1359, 2013.

[110] C. Shui and W. Szeto, "Dynamic green bike repositioning problem–a hybrid rolling horizon artificial bee colony algorithm approach," *Transportation Research Part D: Transport and Environment*, vol. 60, pp. 119–136, 2018.

[111] H. Si, J.-g. Shi, G. Wu, J. Chen, and X. Zhao, "Mapping the bike sharing research published from 2010 to 2018: A scientometric review," *Journal of cleaner production*, vol. 213, pp. 415–427, 2019.

[112] K. Sörensen and N. Vergeylen, "The bike request scheduling problem," in *Computer Aided Systems Theory–EUROCAST 2015: 15th International Conference, Las Palmas de Gran Canaria, Spain, February 8-13, 2015, Revised Selected Papers 15*. Springer, 2015, pp. 294–301.

[113] F. Soriguera *et al.*, "A continuous approximation model for the optimal design of public bike-sharing systems," *Sustainable Cities and Society*, vol. 52, p. 101826, 2020.

[114] S. Tan, Z. Li, and N. Xie, "Dynamic capacitated arc routing problem in e-bike sharing system: A monte carlo tree search approach," *Journal of Advanced Transportation*, vol. 2021, 2021.

[115] Q. Tang, Z. Fu, D. Zhang, H. Guo, and M. Li, "Addressing the bike repositioning problem in bike sharing system: A two-stage stochastic programming model," *Scientific Programming*, vol. 2020, 2020.

[116] C. M. Vallez, M. Castro, and D. Contreras, "Challenges and opportunities in dock-based bike-sharing rebalancing: a systematic review," *Sustainability*, vol. 13, no. 4, p. 1829, 2021.

[117] N. Vergeylen, K. Sörensen, and P. Vansteenwegen, "Large neighborhood search for the bike request scheduling problem," *International Transactions in Operational Research*, vol. 27, no. 6, pp. 2695–2714, 2020.

[118] P. Vogel, "Service network design of bike sharing systems," in *Service Network Design of Bike Sharing Systems*. Springer, 2016, pp. 113–135.

[119] P. Vogel, T. Greiser, and D. C. Mattfeld, "Understanding bike-sharing systems using data mining: Exploring activity patterns," *Procedia-Social and Behavioral Sciences*, vol. 20, pp. 514–523, 2011.

[120] P. Vogel, B. A. Neumann Saavedra, and D. C. Mattfeld, "A hybrid metaheuristic to solve the resource allocation problem in bike sharing systems," in *International workshop on hybrid metaheuristics*. Springer, 2014, pp. 16–29.

[121] M. Wang and X. Zhou, "Bike-sharing systems and congestion: Evidence from us cities," *Journal of transport geography*, vol. 65, pp. 147–154, 2017.

[122] S. Wang, T. He, D. Zhang, Y. Shu, Y. Liu, Y. Gu, C. Liu, H. Lee, and S. H. Son, "Bravo: Improving the rebalancing operation in bike sharing with rebalancing range prediction," *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 2, no. 1, pp. 1–22, 2018.

[123] Y. Wang and W. Szeto, "The dynamic bike repositioning problem with battery electric vehicles and multiple charging technologies," *Transportation Research Part C: Emerging Technologies*, vol. 131, p. 103327, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0968090X21003326

[124] D. J. Watts and S. H. Strogatz, "Collective dynamics of 'small-world'networks," *Nature*, vol. 393, no. 6684, pp. 440–442, 1998.

[125] X. Wu, C. Lyu, Z. Wang, and Z. Liu, "Station-level hourly bike demand prediction for dynamic repositioning in bike sharing systems," in *Smart transportation systems 2019*. Springer, 2019, pp. 19–27.

[126] H. Yang, X. Zhang, L. Zhong, S. Li, X. Zhang, and J. Hu, "Short-term demand forecasting for bike sharing system based on machine learning," in *2019 5th International Conference on Transportation Information and Safety (ICTIS)*. IEEE, 2019, pp. 1295–1300.

[127] Y.-C. Yin, C.-S. Lee, and Y.-P. Wong, "Demand prediction of bicycle sharing systems," 2012, [Online].

[128] P.-S. You, "A two-phase heuristic approach to the bike repositioning problem," *Applied Mathematical Modelling*, vol. 73, pp. 651–667, 2019.

[129] D. Zhang, C. Yu, J. Desai, H. Lau, and S. Srivathsan, "A time-space network flow approach to dynamic repositioning in bicycle sharing systems," *Transportation research part B: methodological*, vol. 103, pp. 188–207, 2017.

[130] J. Zhang, M. Meng, Y. D. Wong, P. Ieromonachou, and D. Z. Wang, "A data-driven dynamic repositioning model in bicycle-sharing systems," *International Journal of Production Economics*, vol. 231, p. 107909, 2021.

[131] X. Zheng, M. Tang, Y. Liu, Z. Xian, and H. H. Zhuo, "Repositioning bikes with carrier vehicles and bike trailers in bike sharing systems," *Applied Sciences*, vol. 11, no. 16, p. 7227, 2021.

[132] P. Zhou, C. Wang, Y. Yang, and X. Wei, "E-sharing: Data-driven online optimization of parking location placement for dockless electric bike sharing," in *2020 IEEE 40th International Conference on Distributed Computing Systems (ICDCS)*, 2020, pp. 474–484.

[133] Y. Zhou, Z. Lin, R. Guan, and J.-B. Sheu, "Dynamic battery swapping strategies for e-bike sharing systems with electric fences," *Available at SSRN 4267760*.

[134] S. Zhu, "Optimal fleet deployment strategy: Model the effect of shared e-bikes on bike-sharing system," *Journal of Advanced Transportation*, vol. 2021, 2021.

APPENDIX A MAINTAINING EQUILIBRIUM BETWEEN THE NUMBER OF BIKES PICKED UP AND DROPPED OFF

This step is crucial as it ensures that the number of bikes in the network remains stable thorugh the simulation.

Let us define four distinct lists that categorize the stations inventory requirements: the lists $\mathcal{S}^{\mathcal{R}+}$ and $\mathcal{S}^{\mathcal{R}+}$ contain alerted stations that require additional regular and electric bikes, respectively; whereas the lists $\mathcal{S}^{\mathcal{R}-}$ and $\mathcal{S}^{\mathcal{E}-}$ contain the stations that require regular and electric bikes to be removed from their inventory. The rebalancing process starts by assigning the stations that raised rebalancing alerts to one list associated to regular bikes ($\mathcal{S}^{\mathcal{R}+}$ or $\mathcal{S}^{\mathcal{R}-}$) and/or to another list associated to electric bikes ($\mathcal{S}^{\mathcal{E}+}$ or $\mathcal{S}^{\mathcal{E}-}$). Stations on all lists are sorted based on their deviation from the target inventory value, ensuring that those with the greatest difference have a higher priority and, therefore, a higher chance of being selected for rebalancing.

The stations are iteratively selected to be rebalanced according to the value of the variables $accumulator^{\mathcal{R}}$ and $accumulator^{\mathcal{E}}$. These variables store the number of regular and electric bikes either exceeding or missing in the network after a sequence of rebalancing operations have occurred – note that negative values signalize missing bikes whereas positive values signalize exceeding bikes in the system. The stations are chosen for rebalancing with the goal of consistently driving the variables $accumulator^{\mathcal{R}}$ and $accumulator^{\mathcal{E}}$ towards zero.

Figure A.1 presents the flowchart of the process of selecting the stations to be rebalanced. The procedure starts by identifying which accumulator ($accumulator^{\mathcal{R}}$ or $accumulator^{\mathcal{E}}$) has the highest absolute value and if its value is positive or not. This step is critical as it dictates from which list ($\mathcal{S}^{\mathcal{R}+}$, $\mathcal{S}^{\mathcal{R}-}$, $\mathcal{S}^{\mathcal{E}+}$ or $\mathcal{S}^{\mathcal{E}-}$) the next station to be rebalanced will be drawn. After the list is selected, the first station of the list, i.e. the station with the highest priority, is selected to be rebalanced regarding both regular and electric bikes. This results in updating $accumulator^{\mathcal{R}}$ and $accumulator^{\mathcal{E}}$, and on removing the selected station from all lists where it is present. That procedure continues until that the number of selected stations reaches the maximum rebalancing capacity. Besides, the procedure is halted whenever it selects, according to the accumulator variables, a list of stations that is empty.

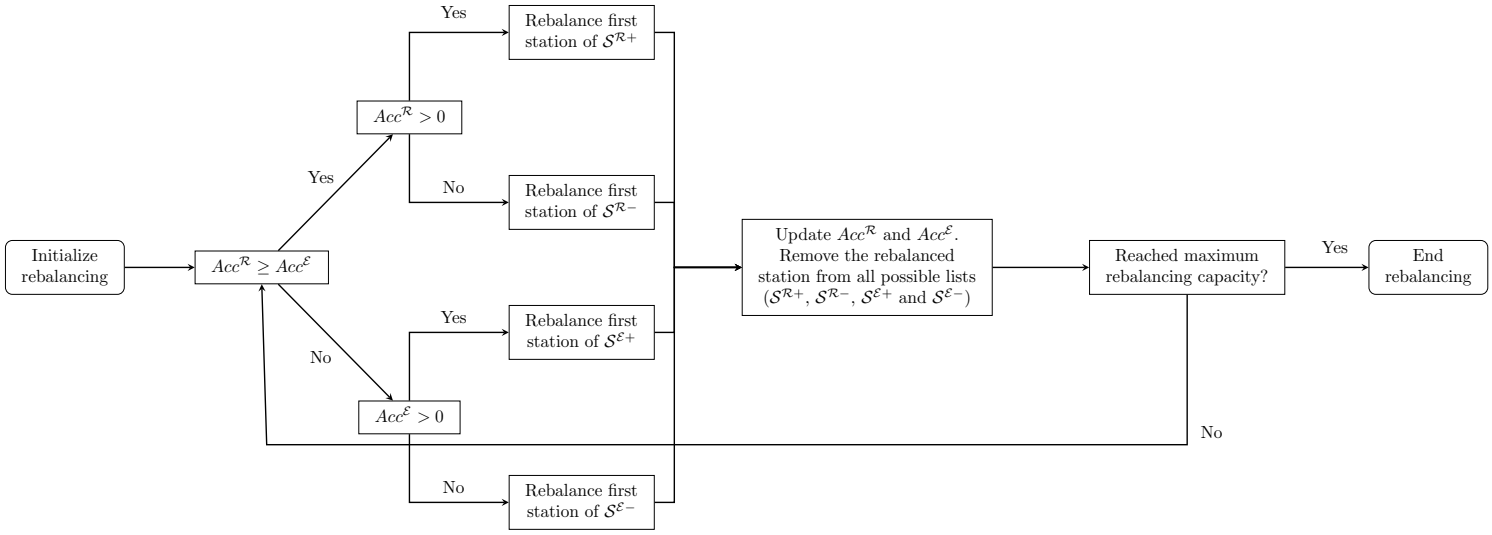

Figure A.1 Flowchart of the selection procedure of stations for rebalancing.

APPENDIX B TRIP PREDICTION AND INVENTORY INTERVAL

The methodology for predicting trips and calculating inventory intervals and target inventories is adapted from the model presented by [53]. The rentals and returns ($f_s^{+,t}$ and $f_s^{-,t}$) are predicted on an hourly basis for each station. The model utilizes a Gradient Boosting Tree, which incorporates weather conditions (temperature and humidity) and temporal information (the day of the week, hour of the day, and holidays) as learning features. In this model, a Singular Value Decomposition (SVD) technique is applied to reduce the dimension of the trip data. This process results in faster predictions, enhancing tractability when dealing with an elevated number of stations. The SVD also elevates the accuracy of the model, indicated by a lower Root Mean Square Error (RMSE), as it effectively eliminates noise and outliers from the trip data. Figure B.1 illustrates the model's pipeline.

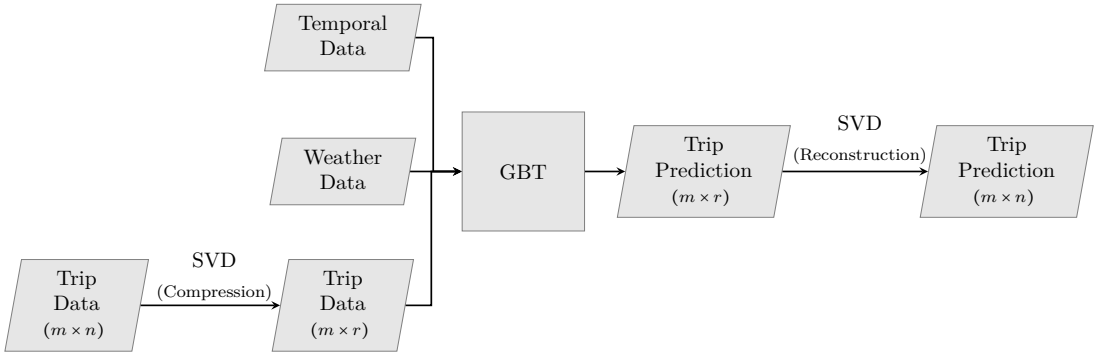

Figure B.1 Pipeline of the predictive model

Based on the predicted rental and return demand, the expected proportion of satisfied trips, known as *service level*, is computed for a given initial inventory. Assuming a station s with initial inventory i and capacity C_s, the rental and the return service levels for a time period $[t, t+\Delta]$ are computed as:

$$SL_s^{+,t}(i) = \frac{\int_t^{t+\Delta} f_s^{+,t}(1 - p_s^t(i,0))dt}{\int_t^{t+\Delta} f_s^{+,t} dt} \tag{B.1}$$

$$SL_s^{-,t}(i) = \frac{\int_t^{t+\Delta} f_s^{-,t}(1 - p_s^t(i,C_s))dt}{\int_t^{t+\Delta} f_s^{-,t} dt}, \tag{B.2}$$

where $f_s^{+,t}$ and $f_s^{-,t}$ represent the rental and return rates, respectively, for station s at time t. Further, $p_s^t(i,0)$, and $p_s^t(i,C_s)$ represent the probability that station s becomes empty and full, respectively, given an initial inventory i at time t.

The *overall service* level is computed as (B.3)

$$SL_s^t(i) = min\{SL_s^{+,t}(i), SL_s^{-,t}(i)\} \tag{B.3}$$

The minimum and maximum service levels, for a station s in time period $[t, t+\Delta]$, can then be computed depending on the initial inventory at time t as follows:

$$SL_s^{min,t} = min_{i\in\{0,\dots,C_s\}}(SL_s^t(i)) \tag{B.4}$$

$$SL_s^{max,t} = max_{i\in\{0,\dots,C_s\}}(SL_s^t(i)). \tag{B.5}$$

A threshold Ω is created to establish an acceptable service level for the time horizon $[t, t+\Delta]$:

$$\Omega_s^t = SL_s^{min,t} + \beta(SL_s^{max,t} - SL_s^{min,t}), \tag{B.6}$$

in which the hyperparameter β controls the proximity of the threshold Ω_s^t to either the minimum service level or the maximum service level. In practice, this hyperparameter influences the gap between the upper and the lower bound of the inventory interval.

The inventory interval for station s for time period $[t, t+\Delta]$ is then defined as

$$\mathcal{I}_s = \{i \in \{0, ..., C_s\} | \mathcal{L} \leq i \leq \mathcal{U}\}, \tag{B.7}$$

where $\mathcal{L} = \min\{i \in \{0, ..., C_s\} | SL_s^t(i) \geq \Omega_s^t\}$, and $\mathcal{U} = \max\{i \in \{0, ..., C_s\} | SL_s^t(i) \geq \Omega_s^t\}$. Finally, the target inventory for station s at time-period $[t, t+\Delta]$ is set to the initial inventory that results in the maximum service level (i.e., $SL_s^{max,t}$).

APPENDIX C WEATHER GENERATOR

The weather generator in Figure 5.3 utilizes normal distributions derived from the temperature and humidity differences between consecutive hours throughout the day, capturing the change in weather conditions over time.

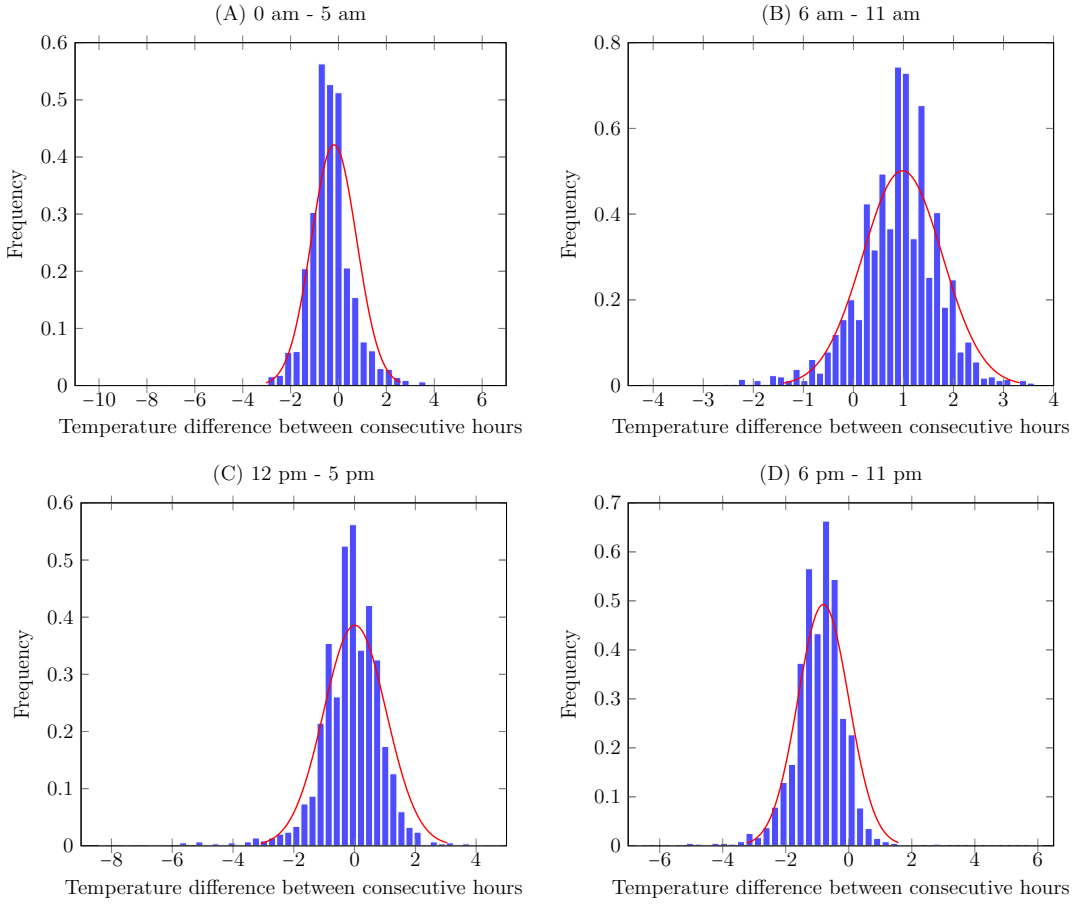

Figure C.1 Distributions of temperature changes between consecutive hours for four time-periods

Histograms of the temperature differences observed between two consecutive hours of 4 periods during the day (0 am – 5 am, 6 am – 11 am, 12 pm – 5 pm, and 6 pm – 11 pm) are depicted in Figure C.1. The overlaid red curves illustrate the normal distributions in which the parameters are computed using the Maximum Likelihood Estimator.

Figure C.1(A) and Figure C.1(C) display narrower distributions, suggesting less variability in temperature change, whereas Figure C.1(B) and Figure C.1(D) exhibit wider spreads,

indicating greater fluctuation. The distributions for humidity differences are obtained using the same approach.

APPENDIX D EXPERIMENTAL RESULTS ON BIXI CLUSTER

The daily lost demand of models DROB-LD and DROB-I on the considered BIXI cluster are detailed in Table D.1.

Table D.1 Lost demand for DROB-LD and DROB-I on BIXI cluster

Day	Model	Opt.Time (min)	MIP Gap (%)	Lost demand (%)		
				Rental	Return	Total
Day 1	DROB-LD	0.00	0.00	0.85	7.52	4.40
	DROB-I	0.01	0.00	0.00	2.26	1.20
Day 2	DROB-LD	0.66	0.00	0.50	0.55	0.52
	DROB-I	0.41	0.00	0.50	0.55	0.52
Day 3	DROB-LD	5.41	0.00	0.73	0.98	0.85
	DROB-I	0.01	0.70	0.73	1.96	1.33
Day 4	DROB-LD	0.06	0.00	2.96	1.51	2.27
	DROB-I	0.01	0.00	2.09	1.51	1.81
Day 5	DROB-LD	0.02	0.00	1.74	2.33	2.01
	DROB-I	0.01	0.00	1.90	4.83	3.27
Day 6	DROB-LD	0.01	0.00	0.16	0.18	0.17
	DROB-I	0.01	0.00	0.00	1.41	0.66
Day 7	DROB-LD	0.04	0.00	0.35	2.92	1.56
	DROB-I	0.01	0.00	0.17	3.31	1.65
Day 8	DROB-LD	0.01	0.00	2.66	0.65	1.76
	DROB-I	0.01	0.00	2.66	2.39	2.54
Day 9	DROB-LD	0.09	0.00	1.75	0.80	1.30
	DROB-I	0.01	0.00	1.92	1.59	1.77
Day 10	DROB-LD	0.02	0.00	2.90	2.49	2.71
	DROB-I	0.01	0.00	2.90	2.49	2.71
Day 11	DROB-LD	0.01	0.00	1.71	0.73	1.23
	DROB-I	0.01	0.00	1.54	1.63	1.58
Day 12	DROB-LD	0.02	0.00	0.65	1.65	1.12
	DROB-I	0.01	0.00	1.13	0.74	0.95
Day 13	DROB-LD	0.13	0.00	3.58	1.42	2.52
	DROB-I	0.03	0.00	2.21	0.18	1.22
Day 14	DROB-LD	0.01	0.00	2.62	1.32	1.99
	DROB-I	0.01	0.00	2.62	1.81	2.23
Day 15	DROB-LD	0.01	0.00	0.85	2.21	1.50
	DROB-I	0.01	0.00	0.68	0.37	0.53
Day 16	DROB-LD	0.02	0.00	0.53	3.03	1.81
	DROB-I	0.02	0.00	0.00	3.03	1.55
Day 17	DROB-LD	0.11	0.00	2.56	1.12	1.84
	DROB-I	0.01	0.00	2.28	0.00	1.13
Day 18	DROB-LD	0.36	0.00	2.80	0.76	1.79
	DROB-I	0.03	0.00	2.24	3.82	3.02
Day 19	DROB-LD	0.01	0.00	2.34	0.20	1.29
	DROB-I	0.01	0.00	0.78	1.62	1.19
Day 20	DROB-LD	0.01	0.00	0.40	1.08	0.72
	DROB-I	0.01	0.00	1.00	0.43	0.72