Dynamic Shared Bicycle Repositioning with Flexible Time Horizon

T.W.K. (Tom) Bruinink

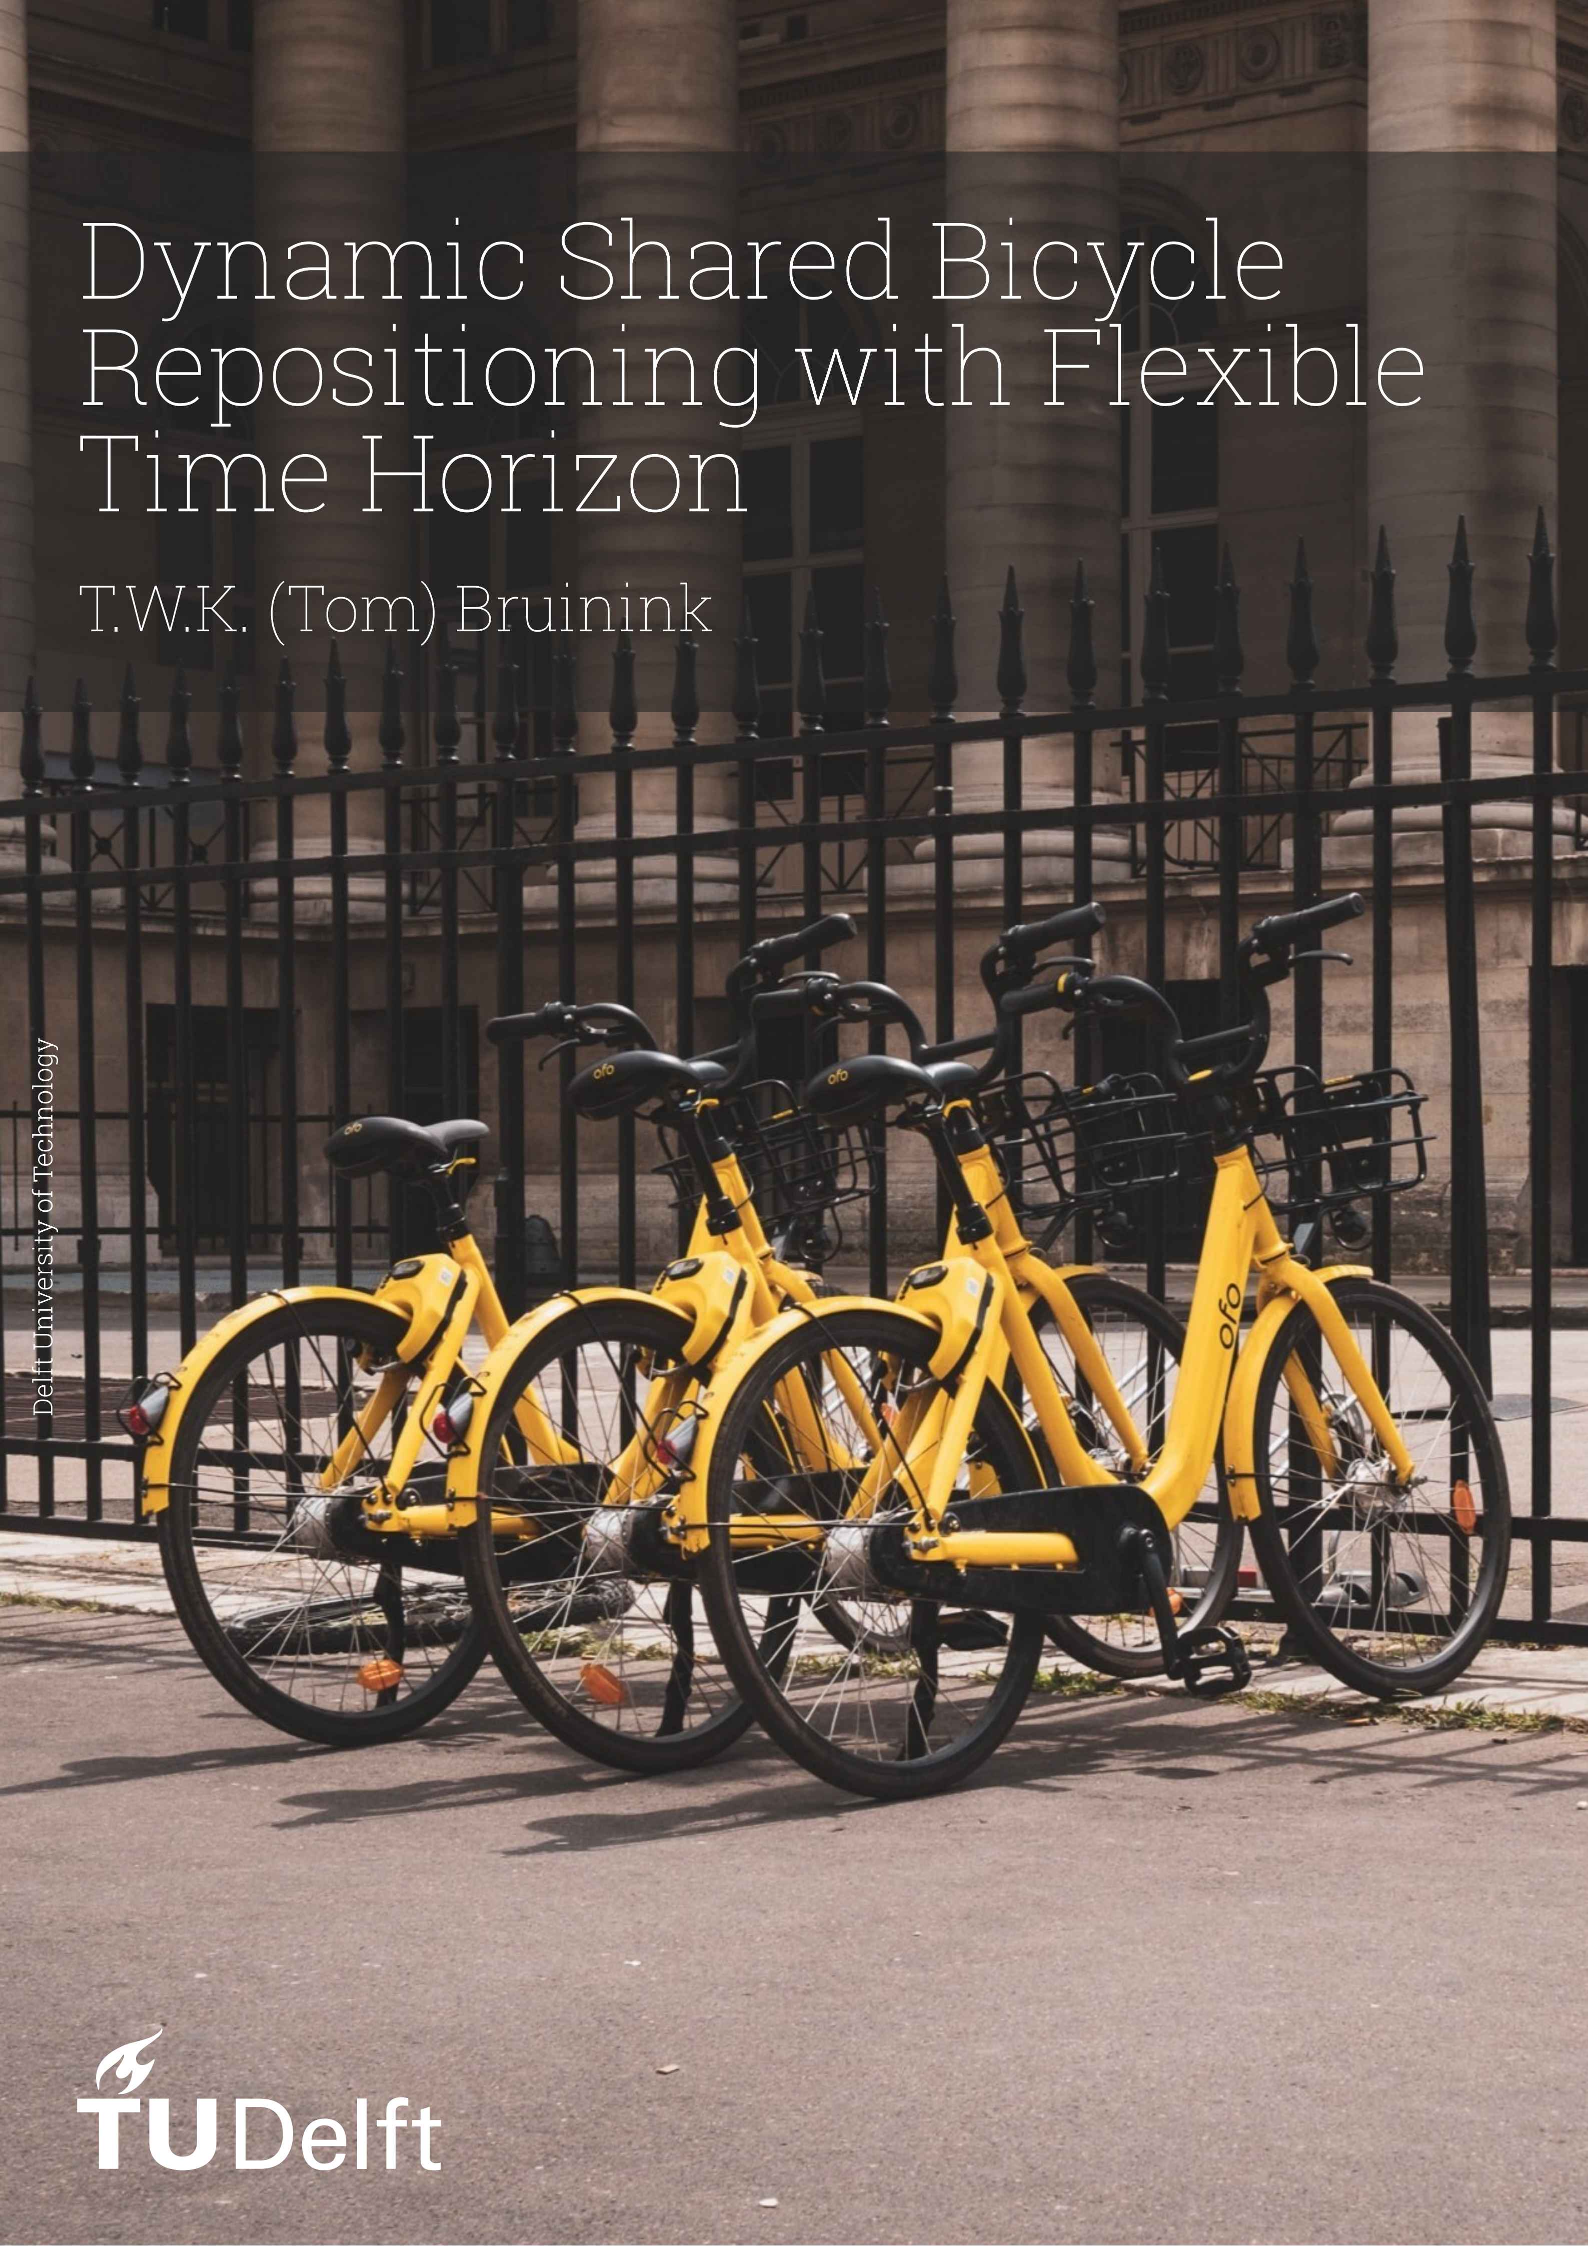

Delft University of Technology

Dynamic Shared Bicycle Repositioning with Flexible Time Horizon

by

T.W.K. (Tom) Bruinink

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Monday September 8, 2025 at 15:00 PM.

Student number:	6075673	
Project duration:	March, 2025 – August, 2025	
Thesis committee:	S. Sharif Azadeh	Chair (TU Delft)
	M. Y. Maknoon	Supervisor (TU Delft)
	R. Artime Torres	Company supervisor (Nommon)
	M. B. Duinkerken	Additional examiner (TU Delft)

An electronic version of this thesis is available at http://repository.tudelft.nl/.

TU Delft

Acknowledgements

This thesis marks the completion of my MSc in Transport, Infrastructure, and Logistics at TU Delft. Throughout this journey, I have had the privilege of working with and learning from many inspiring individuals. Without their support, this thesis would not have been possible. I am grateful to have explored a topic that aligns with my interests in transportation and mobility, while also contributing to sustainability and accessibility with the ultimate aim of having a positive impact on society.

First and foremost, I would like to express my sincere gratitude to Nommon Solutions and Technologies for the opportunity to conduct this research project within their company. The team's passion for mobility has been truly inspiring and my time there has allowed me to grow both professionally and personally. I also had the opportunity to improve my Spanish and gain a deeper appreciation of Spanish culture. I would like to especially thank Ruben Artime Torres for his daily supervision, his belief in me from the beginning, and the freedom he gave me to explore the research in the direction I saw fit. My sincere thanks also go to Ines Peirats for sharing her expertise on shared bicycle systems and always answering my many questions, and to Pablo Ruiz for providing demand forecasts and brainstorming about the simulation model. Lastly, I am thankful to all my Nommon colleagues, you made me feel at home from day one, involved me in many activities, and patiently listened to my broken Spanish.

In academia, I am grateful to my supervisors, Shadi Sharif Azadeh and Yousef Maknoon, for their invaluable feedback, guidance, and continued support. Your insights helped shape the direction of my research and challenged me to find better solutions. I would also like to thank Mark Duinkerken for being part of my thesis examination committee, your time and contribution are greatly appreciated.

Lastly, I would like to thank my family and friends for their unlimited support throughout my academic journey. Your encouragement to broaden my horizon by pursuing a second master's degree and an internship abroad has been very valuable. A special word of thanks goes to my father, who read and reread my thesis with care, offering many valuable textual suggestions.

As this academic chapter comes to a close, I look forward to the future with enthusiasm. My academic journey across different universities has allowed me to explore a wide range of topics. These experiences have inspired me to apply the knowledge that I have gained to contribute to improving society.

T.W.K. (Tom) Bruinink
Madrid, August 2025

Summary

Shared bicycle systems have become an important element of sustainable urban mobility, helping to reduce emissions and congestion while improving first- and last-mile connectivity. However, these systems often suffer from imbalances between bicycle supply and demand, especially during peak hours, reducing user satisfaction and system reliability.

This thesis addresses these challenges by proposing a dynamic, time-adaptive repositioning strategy for continuously operating shared bicycle systems. The strategy integrates station selection, bicycle quantity decisions, and routing. It adapts the planning horizon based on the time of day using shorter routes during peak periods and longer routes during off-peak hours and at night. User satisfaction is captured through a no-service penalty, while operational efficiency is ensured through time-dependent constraints on route duration. The model is further extended to account for mixed bicycle fleets and applies vehicle-aware repositioning to improve coordination and reduce computational complexity.

Real-world case studies in Zaragoza and Valladolid show that the dynamic strategy significantly outperforms both static and no-repositioning methods. In Zaragoza, it enabled approximately 400 additional satisfied trips per day, while in Valladolid it achieved exceeding a satisfaction level of 95% despite a high demand for electric bicycles. The dynamic approach also proved robust to a 50% increase in demand for both systems, with only minimal reductions in satisfaction levels.

In summary, this thesis demonstrates that a time-adaptive and integrated repositioning strategy can greatly enhance service levels and user satisfaction in shared bicycle systems. The approach supports broader adoption and strengthens the role of shared bicycles in sustainable urban transportation. Future research could build on these findings by targeting bottlenecks with customized strategies, managing broken bicycles, which reduce dock availability, and accounting for solar-powered stations, without a recharge option, where low battery bicycles occupy dock space.

Contents

1
Introduction

Shared bicycle systems have become a key element of sustainable urban transportation by reducing greenhouse gas emissions, travel time, and traffic congestion, while improving connectivity to other modes of transit by addressing the first-mile/last-mile challenge (DeMaio, 2009). Over the last two decades, the number of shared bicycle systems has expanded rapidly, with more than 2000 systems and nearly 10 million bicycles operating worldwide (O'Brien, 2025).

However, the rapid expansion of shared bicycle systems has also introduced new challenges. With an increasing number of stations and users, maintaining a balance between supply and demand has become increasingly complex. A major issue is the imbalance in bicycle availability between different areas throughout the day. During the morning peak, a large flow of bicycles moves from residential areas to business districts, while the opposite flow occurs in the evening. This natural movement causes some stations to become empty, leaving users unable to pick up a bicycle, while others become overcrowded, making it impossible to return one. This imbalance reduces the system's efficiency and user satisfaction, ultimately limiting the growth of users and the overall impact of the system.

To maintain system stability, operators use repositioning vehicles to redistribute bicycles across stations. Traditionally, repositioning has been reactive, with operators responding to imbalances only after stations become empty or full. This strategy often leads to delays, inefficient routing, and higher costs. As a result, imbalances persist longer, making the service less reliable to users. To overcome these challenges, there is a growing shift toward proactive repositioning, where operators anticipate future imbalances using demand predictions. By taking proactive action, they can mitigate potential shortages or surpluses before they disrupt the system, improving service reliability and operational efficiency.

The most common strategy to address the bicycle repositioning problem is a static approach, where bicycles are redistributed during off-peak hours, mainly at night (Alvarez-Valdes et al., 2016; Chemla et al., 2013). In this approach, repositioning routes are planned to prepare the system for the expected demand of the following day. However, this method struggles to respond to real-time system dynamics and heavily relies on demand predictions, which are often uncertain. Furthermore, solving the resulting Mixed Integer Programming (MIP) formulation is computationally too complex for large-scale systems, requiring the use of heuristics and approximations. Additionally, preventing shortages during the day would require an excessively large fleet of bicycles and number of docks, making this approach inefficient for dynamic environments.

An alternative strategy is the dynamic approach, in which repositioning occurs during operation, based on the real-time state of the system and short-term demand forecasts. Repositioning vehicles are deployed on short routes that serve only the most critical stations. This method offers several advantages: It can respond to fluctuations in demand, depends less on demand forecasts, and has a low computational cost, as typically only one vehicle is optimized at a time, specifically the one that has just completed its previous route.

The dynamic approach to bicycle repositioning can take several forms, each varying in how decisions are made and actions are planned. One variant involves assigning the next repositioning action, such as picking up or delivering bicycles, to a vehicle immediately after it completes its previous task. Although this method enables fast and real-time decision-making, it falls short in terms of routing efficiency, as each action is treated independently without the possibility of clustering multiple actions within a short time frame. Another variant follows a two-step structure, where the system first determines either the optimal number of bicycles to reposition or the route to follow, and subsequently solves the remaining part of the problem based on that initial decision. Although this method offers a structured framework, it often lacks flexibility, as the quality of the final outcome strongly depends on the first-stage choices. A more comprehensive variant integrates repositioning and routing decisions into a single model. This integrated approach simultaneously determines which stations to visit, how many bicycles to move, and the optimal sequence of actions. By addressing all subproblems together, it allows vehicles to receive optimized routes as soon as they become available for reassignment.

Dynamic bicycle repositioning can be planned and updated over different time horizons. The rolling-horizon method splits the dynamic problem into a series of smaller static subproblems by optimizing decisions within fixed time steps. In contrast, time-space network flow models represent the problem in continuous time, enabling the capture of temporal dynamics and coordination among multiple vehicles, but at the cost of increased computational complexity. To support continuous operation, a time-adaptive framework, considered a variant of the rolling-horizon approach, generates repositioning routes in real time while dynamically adjusting the planning horizon based on the time of day. This introduces temporal flexibility: longer routes during off-peak or overnight periods improve operational efficiency, while shorter routes during peak hours respond quickly to rapid demand fluctuations, enhancing service quality.

An extension of the dynamic approach addresses the growing prevalence of mixed bicycle fleets, particularly with the increasing integration of electric bicycles (Ghamami & Shojaei, 2018). Incorporating electric bicycles into the repositioning strategy enhances system accessibility, enables longer trips, and reduces physical effort, making the service more appealing and practical to a wider range of users (Galatoulas et al., 2020).

To manage computational complexity, especially in large-scale systems, various strategies have been proposed. A straightforward and commonly used approach is to focus solely on rebalancing imbalanced stations, thereby limiting the number of stations considered. In addition to this, two location-based strategies have been suggested. The first involves clustering stations into fixed service zones, reducing the problem size by allowing vehicles to operate within specific areas or enabling the problem to be solved in multiple stages. However, this comes at the cost of reduced vehicle flexibility. An alternative and more adaptive approach, known as vehicle-aware repositioning, accounts for vehicle end locations to guide potential future assignments. This forward-looking method narrows the solution space without imposing rigid spatial boundaries, thereby improving fleet coordination and scalability.

1.1. Methodology

This thesis proposes an integrated dynamic repositioning model that simultaneously optimizes station selection, bicycle movements, and routing in real time. The model incorporates a flexible planning horizon, which is shorter during peak hours and longer during off-peak periods, to balance service quality and operational efficiency. The model will be extended to include a mixed bicycle fleet. This extension increases the complexity of the problem, since repositioning decisions must account for both bicycle types. To ensure scalability and coordination in multi-vehicle operations, the model uses vehicle-aware repositioning, taking into account the expected end location of all vehicle to guide future task assignments.

Building on these elements, this thesis makes four main contributions. First, it defines operational efficiency and user satisfaction measures specifically for continuously operating shared bicycle systems, where repositioning vehicles are in constant operation and therefore minimizing operating time is not a relevant objective. Second, it introduces a time-adaptive dynamic repositioning strategy that accounts for the changing conditions throughout the day. This strategy integrates station selection, bicycle quantity decisions, and efficient route planning within a time-dependent planning horizon, with the primary objective of maximizing user satisfaction while using a flexible maximum route duration.

Third, the model is extended to address practical complexities, including the presence of a mixed bicycle fleet and the need for more coordination between operating vehicles. To address this, we introduce vehicle-aware repositioning, a forward-looking approach that takes into account the expected end locations of other vehicles into the decision-making process. This reduces computational complexity and improves operational efficiency, without creating fixed service areas. Finally, based on a realistic case study, the results offer valuable insights into the growing importance of dynamic bicycle repositioning for sustaining service quality and ensuring system scalability.

1.2. Structure

This thesis is structured as follows. Chapter 2 reviews the existing literature, while Chapter 3 outlines the research questions. Chapter 4 introduces the problem and Chapter 5 presents the model. The model extensions are discussed in Chapter 6. Chapter 7 explains the simulation method and Chapter 8 discusses the case studies. Chapter 9 presents the results. Finally, Chapter 10 and Chapter 11 summarize the main conclusions and propose directions for future research.

2

Literature review

This literature review begins with an explanation of the static approach to bicycle repositioning, which was the original method but has notable limitations. We then turn to the dynamic approach, which addresses many of these shortcomings and forms the main focus of our review. Within this dynamic framework, we examine various methodological variants, followed by a discussion of possible extensions aimed at further improving model performance. An overview of the main characteristics of the literature reviewed in this chapter is provided in Table 2.1. The table classifies studies according to their repositioning strategy, optimization method, planning horizon, demand prediction type, fleet composition, coordination level, and primary objective. Only papers explicitly addressing the bicycle repositioning problem are included; studies focusing exclusively on the station inventory problem, without routing or relocation considerations are excluded due to their fundamentally different scope and modeling assumptions.

Table 2.1: Classification of literature on bicycle repositioning approaches

Reference	Strategy	Optimization method	Time horizon	Demand predictions	Fleet type	Coordination	Main objective
Static approaches							
Chemla et al. (2013)	Integrated	Exact/Heuristic	Fixed	-	Single	Clustering	Cost optimization
Rainer-Harbach et al. (2013)	Two-step	Heuristic	Fixed	-	Single	No coordination	User satisfaction
Alvarez-Valdes et al. (2016)	Two-step	Exact	Fixed	Long-term	Single	No coordination	User satisfaction
Li et al. (2016)	Integrated	Heuristic	Fixed	Long-term	Mixed	No coordination	Hybrid
Schuijbroek et al. (2017)	Integrated	Heuristic	Fixed	Long-term	Single	Clustering	Cost optimization
Zhu (2021)	Two-step	Heuristic	Fixed	-	Mixed	No coordination	Cost optimization
Dynamic approaches							
Caggiani and Ottomanelli (2013)	Integrated	Simulation-based	Rolling horizon	Short-term	Single	No coordination	Hybrid
Kloimüllner et al. (2014)	Integrated	Heuristic	Rolling horizon	Short- & Long-term	Single	No coordination	User satisfaction
Pfrommer et al. (2014)	Two-step	Heuristic	Rolling horizon	Short-term	Single	Vehicle-aware	User satisfaction
Regue and Recker (2014)	Two-step	Exact	Rolling horizon	Short-term	Single	No coordination	Hybrid
O'Mahony and Shmoys (2015)	Integrated	Exact/Heuristic	Time-adaptive	Short- & Long-term	Single	Clustering	User satisfaction
Brinkmann et al. (2015)	Action-based	Simulation-based	Rolling horizon	-	Single	No coordination	User satisfaction
Zhang et al. (2017)	Two-step	Exact/Heuristic	Time-space	Short-term	Single	No coordination	Hybrid
Shui and Szeto (2018)	Two-step	Heuristic	Rolling horizon	Short-term	Single	No coordination	Hybrid
Caggiani et al. (2018)	Integrated	Simulation-based	Rolling horizon	Short-term	Single	Clustering	User satisfaction
Brinkmann et al. (2019)	Action-based	Simulation-based	Time-adaptive	Short- & Long-term	Single	No coordination	User satisfaction
Legros (2019)	Action-based	Exact	Rolling horizon	Short-term	Single	Clustering	User satisfaction
Jiménez-Meroño and Soriguera (2024)	Action-based	Heuristic	Rolling horizon	Short-term	Single	Vehicle-aware	User satisfaction
This study	Integrated	Exact	Time-adaptive	Short-term	Mixed	Vehicle-aware	User satisfaction

2.1. Static approach

The static repositioning approach treats the system as "closed" for the purpose of repositioning, which means that bicycle movements are not considered during operation (Alvarez-Valdes et al., 2016; Rainer-Harbach et al., 2013). Repositioning typically occurs at night when bicycle use is minimal and demand is predictable.

The aim is to optimize the distribution of bicycles at the start of each day by predicting future demand at each station. Raviv and Kolka (2013) propose a method that uses an inventory model to track bicycle availability over time, accounting for varying demand patterns between stations. By incorporating expected rental and return patterns, the model helps to maintain inventory balance throughout the day.

As a result, repositioning efforts during the day are minimized, leading to significant cost savings, as repositioning during the day is more expensive. This approach is further improved by Schuijbroek et al. (2017), who enhances the model by using Markov processes to establish lower and upper bounds for the optimal station inventory, allowing more flexibility.

To achieve these optimal inventory levels, bicycles need to be relocated between stations. This problem can be viewed as a one-commodity pickup and delivery problem (Hernández-Pérez & Salazar-González, 2004; Salazar-González & Santos-Hernández, 2015), where bicycles are the commodity being transported between stations, similar to a vehicle routing problem (VRP). In this approach, all routes are planned simultaneously, typically visiting all stations to achieve the desired bicycle distribution while minimizing routing costs (Alvarez-Valdes et al., 2016). This problem is often formulated as a mixed integer programming (MIP) optimization model (Chemla et al., 2013; Salazar-González & Santos-Hernández, 2015), which is an exact method for finding a solution. However, since the problem has been proven to be NP-hard (Chemla et al., 2013), which means that its complexity increases non-polynomially with the size of the system, solving it exactly for real-world problems becomes computationally expensive. As a result, heuristics and approximations are often employed to find feasible solutions (Jiménez-Meroño & Soriguera, 2024).

Despite its advantages, the static repositioning approach has significant limitations. It cannot respond to real-time changes in bicycle demand or availability during operational hours. Since the number of available bicycles is most relevant for prediction horizons of only a few hours, the model is less effective for longer periods of time (Gast et al., 2015). Furthermore, to prevent imbalances throughout the day, an unreasonably large bicycle fleet would be required, making this static approach inefficient for dynamic systems (Jiménez-Meroño & Soriguera, 2024). Therefore, while static models offer an initial solution, they struggle to accommodate the dynamic nature of shared bicycle systems.

2.2. Dynamic approach

The goal of the dynamic approach to bicycle repositioning is to maintain balanced bicycle availability across stations in real time, ensuring high user satisfaction by minimizing empty and full stations during operational hours. Unlike static repositioning, which is heavily dependent on long-term demand forecasts and preplanned routes, the dynamic approach leverages real-time station data and short-term demand predictions to make adaptive and responsive decisions throughout the day. A key feature of this method is that new repositioning routes are generated only after the previous route is completed. This sequential decision-making process significantly reduces computational complexity, as only one route is constructed at a time, rather than solving an entire network of repositioning actions simultaneously (Jiménez-Meroño & Soriguera, 2024). The primary objective of dynamic repositioning is to minimize unsatisfied demand by proactively adjusting station inventories before they become empty or full. While factors such as minimizing travel time or aligning station inventories with target fill levels are still considered, they are generally secondary to service quality (Jiménez-Meroño & Soriguera, 2024). Studies have shown that even minimal repositioning efforts can lead to substantial improvements in service levels (Legros, 2019). Interestingly, the optimal inventory level for each station often varies throughout the day based on demand fluctuations but tends to be independent of vehicle capacity (Legros, 2019). To support these decisions, various prioritization policies have been developed to select which stations to visit. These include targeting stations with high current or future unmet demand, large deviations from desired inventory levels, high movement frequency, or strategic proximity to other stations (Legros, 2019). In contrast, stations with consistently low demand are deprioritized to avoid inefficient repositioning operations (Jiménez-Meroño & Soriguera, 2024).

There are different strategies within the dynamic approach to bicycle repositioning, each varying in how decisions are made and how repositioning actions are planned and executed. One strategy focuses on assigning the next repositioning task, typically a single pickup or delivery action, to a vehicle immediately after it completes its previous task. This action-based strategy significantly reduces computational complexity by avoiding full route planning and instead focuses on individual high-priority actions based on real-time station data and short-term demand forecasts (Jiménez-Meroño & Soriguera, 2024; Legros, 2019). This allows for quick responsiveness to demand fluctuations and ensures that the most critical stations are served in a timely manner. For example, Jiménez-Meroño and Soriguera (2024) and Legros (2019) model the problem as a pairwise task assignment, where a repositioning vehicle is continuously assigned to a single action. Likewise, Brinkmann et al. (2015) introduce a framework that distinguishes between short- and long-term strategies. Their short-term approach prioritizes nearby stations with immediate violation risks, while the long-term approach targets the most imbalanced stations overall. Notably, their model shifts the objective from the quantity of unmet demand to the duration of due date violations, penalizing the time that stations remain full or empty after user requests. However, a common limitation of these action-based strategies is that decisions are made sequentially, considering only one step at a time. This can lead to inefficiencies as opportunities to combine actions or optimize across multiple stations are missed. For example, visiting two nearby unbalanced stations together or allowing simultaneous pickups and deliveries might produce better results than handling each action in isolation (Legros, 2019).

Another strategy is the two-step method, which separates decision-making into sequential optimization phases. In the version proposed by Pfrommer et al. (2014), the first step involves identifying promising repositioning routes based on their expected utility. In the second step, the optimal number of bicycles to relocate along each route is determined, after which the route with the highest overall utility is selected for execution. A slightly different two-step approach is described by Regue and Recker (2014), where the process begins by determining the optimal inventory levels at the stations. Based on these targets, a route is constructed for a single repositioning vehicle. To ensure computational feasibility, only stations with significant inventory imbalances and within a threshold distance from the vehicle's current location are considered. This restriction bounds the size of the problem and enables the use of traditional solvers. The method also allows pickups or drop-offs at balanced "buffer" stations to avoid situations where the vehicle becomes empty or overfull, thereby increasing operational flexibility. A common limitation of two-step approaches lies in the strong influence that the first-stage decision has on the overall outcome. Although fixing part of the solution early can improve computational efficiency and structure the problem, it also reduces flexibility in the second stage, potentially leading to suboptimal results.

The integrated approach to dynamic bicycle repositioning combines the decisions of determining which stations to visit, the optimal number of bicycles to relocate, and the most efficient routing (Kloimüllner et al., 2014). By addressing all of these subproblems together in a single model, the integrated approach enables vehicles to receive optimized repositioning routes. This approach ensures that the repositioning process is highly responsive to real-time demand fluctuations, ultimately improving both operational efficiency and user satisfaction.

An important dimension in dynamic bicycle repositioning is the time horizon over which decisions are planned and updated. Two common formulations are the rolling-horizon and the time–space network approaches, which can be applied to different repositioning strategies regardless of whether they are integrated, two-step, or action-based. The rolling-horizon method divides the operational period into a sequence of fixed-duration intervals, each treated as a static subproblem (Shui & Szeto, 2018). Within each stage, real-time station statuses and short-term demand forecasts are assumed to be known in advance, effectively making the problem static for that interval. Once the repositioning routes are executed, the system is updated with new inventory levels and demand estimates, and the next stage begins. This approach maintains adaptability to changing conditions while simplifying the optimization process at each step. A common implementation of this method is to divide the day into five-minute intervals (Caggiani & Ottomanelli, 2013; Caggiani et al., 2018; Shui & Szeto, 2020). This short interval lengths provides sufficient opportunities to react to fluctuations in demand while keeping the computational effort manageable. However, assumptions about how repositioning actions align with these time steps vary. In the work of Caggiani et al. (2018), a single task must be completed in a single five-minute step, ensuring a clear separation between consecutive tasks. In contrast, Shui and Szeto (2018) allow repositioning operations to span multiple time steps, offering greater flexibility and realism, particularly for longer travel distances or more complex routes. The time–space network approach instead models operations on a continuous-time basis, providing a more detailed representation of system dynamics. This method, as explored by Zhang et al. (2017), facilitates real-time decision-making by capturing the evolving nature of both the repositioning process and station demand. A key advantage of this approach is its ability to take into account the actual station status when the repositioning vehicle arrives, resulting in more accurate repositioning actions. However, flexibility in modeling comes with a significant increase in computational complexity. To manage this, Zhang et al. (2017) proposed a decomposition into a two-stage optimization model, reformulating the original nonlinear problem into a mixed-integer problem. Although this makes the problem more tractable, such a staged approach inherently leads to suboptimal solutions, as decisions in the first stage do not fully anticipate outcomes of the second. As a result, this approach has seen limited application in the literature and remains impractical for large-scale, real-world systems.

The dynamic approach can be extended in several ways to further enhance system performance. One extension is the time-adaptive approach, considered as a variant of the rolling-horizon method, which accounts for the heterogeneous demand patterns throughout the day, recognizing that the operator's priorities typically align with these varying demand levels. O'Mahony and Shmoys (2015) propose using a clustering model during the day and a routing model overnight, acknowledging that different periods require tailored strategies. Brinkmann et al. (2019) propose a policy that simulates future demand over a predefined horizon, with decision-making horizons that vary throughout the day to capture heterogeneous demand patterns. These time-dependent look-ahead horizons are autonomously set using value function approximation, enabling the system to anticipate peak-hour commuter demand without extending simulations unnecessarily. From an operational perspective, their model assumes that repositioning vehicles operate continuously throughout the day and that costs related to routing and drivers are already incurred. As a result, transportation costs are excluded from the optimization, allowing the model to focus primarily on maintaining high service levels.

Another valuable extension is the consideration of a mixed fleet, involving multiple types of bicycles. Shui and Szeto (2020) point out that most existing dynamic repositioning models assume homogeneous fleets, which limits their applicability to real-world systems that often offer a variety of bicycle types, such as electric and mechanical bicycles. The model proposed by Li et al. (2016) introduces this heterogeneity by considering bicycles with varying characteristics, such as multi-seat configurations. Zhu (2021) studies how the integration of electric bicycles into a mechanical fleet impacts system performance. This introduces new operational challenges: A surplus of one type of bicycle cannot necessarily compensate for a shortage of another, and different types of bicycles may occupy varying amounts of space in both repositioning vehicles and docking stations. To address this, Martins Silva et al. (2023) determine target inventory levels for each bicycle type and examine how the willingness of users to substitute their preferred bicycle type influences system performance. These complexities highlight the need for more refined loading and unloading strategies, as well as more advanced decision-making processes to ensure that the correct type of bicycle is delivered to meet the specific user demand.

A final extension aimed at reducing computational complexity involves considering vehicle locations in the decision-making process. Traditionally, many studies address complexity by clustering stations into distinct service areas, each assigned to a repositioning vehicle. For example, Chemla et al. (2013), Schuijbroek et al. (2017) and Legros (2019) divide the city into exclusive zones to avoid coordination issues and allow independent vehicle routing. Similarly, O'Mahony and Shmoys (2015) propose clustering stations during peak hours to manage high demand and system fluctuations. Clustering helps reduce the problem size and avoids long truck routes, which are unreliable due to extensive traffic and fast-changing demand. Stations with similar rush-hour behavior are grouped into clusters (Caggiani et al., 2018), each assigned a target service level based on typical usage patterns. Although clustering offers clear computational advantages, it can limit flexibility, particularly in dynamic environments. To address this, a more adaptive strategy, vehicle-aware repositioning, directly incorporates the locations of other repositioning vehicles into the model. Pfrommer et al. (2014) and Jiménez-Meroño and Soriguera (2024), for example, explicitly consider the end locations of all other repositioning vehicles in the system when assigning a new task. This approach preserves some of the efficiency of clustering, while enhancing flexibility by enabling decisions that account for both current and anticipated vehicle positions.

2.3. Our contributions

An analysis of the current literature reveals that most studies on bicycle repositioning focus on static repositioning during the night or dynamic repositioning during the day, often overlooking the changing conditions throughout the day. In contrast, we consider systems that operate continuously using time-dependent planning horizons. Unlike traditional approaches that prioritize cost minimization, our model focuses exclusively on maximizing user satisfaction and incorporates a flexible maximum route duration. We propose an integrated model that simultaneously determines which stations to visit, how many bicycles to relocate, and the most efficient route to follow. Although most existing research assumes a homogeneous bicycle fleet, the introduction of electric bicycles adds complexity that requires additional operational constraints. To address this, we extend our model to accommodate a mixed fleet. Finally, to reduce computational complexity and enhance coordination among vehicles, we apply vehicle-aware repositioning, which introduces restrictions based on the end locations of other vehicles. This method excludes locations near the end locations of other vehicles when assigning the next task.

3 Research Questions

The objective of this research is to enhance the repositioning of shared bicycles by means of a dynamic, flexible time-horizon approach that optimizes user satisfaction while ensuring operational efficiency. By doing so, the system can attract more users and increase the overall impact of shared bicycles. To achieve this, the following main research questions need to be answered:

> How can operational efficiency and user satisfaction be measured in continuously operating shared bicycle systems?
>
> How can a time-adaptive dynamic bicycle repositioning model, integrating real-time station data and demand forecasts, optimize operational efficiency and user satisfaction?
>
> How can a dynamic bicycle repositioning model be extended for a mixed bicycle fleet?

The main research questions are jointly answered by the following sub-questions:

1. How can a repositioning strategy be used for continuously operating shared bicycle systems?
2. How can vehicle-aware repositioning enhance operational efficiency and reduce computational complexity?

Chapter 4 addresses the first sub-question. It explains how a continuously operating shared bicycle system can be repositioned and identifies the inputs required for a repositioning strategy. The first research question is tackled at the beginning of Chapter 5, where operational efficiency and user satisfaction are defined for continuously operating systems. In the remainder of that chapter the second research question is answered with a description of the development of the time-adaptive dynamic repositioning model and an explanation of how real-time station data and demand forecasts are used. The third research question is addressed in Chapter 6, where it is explained how the model is extended to manage a mixed fleet of mechanical and electric bicycles. This chapter also addresses the second sub-question by introducing vehicle-aware repositioning, a method that incorporates the expected end locations of other vehicles to improve coordination and reduce computational complexity.

Finally, Chapter 7 presents the simulation tool used to evaluate the proposed strategies, while in Chapter 8 the application to realistic case studies is discussed. The resulting analysis in Chapter 9 shows the impact of these choices on operational efficiency and user satisfaction, highlighting the key trade-offs.

4 Problem definition

To evaluate the trade-off between operational efficiency and user satisfaction, we consider a daily operation of a shared bicycle system. The system must function across different time periods, overnight, peak, and off-peak hours, each characterized by different demand patterns. Throughout the day, multiple repositioning vehicles continuously relocate bicycles to maintain a balanced inventory across stations, based on the current number of bicycles at each station and the expected demand in the upcoming period. Each driver completes one task at a time, usually involving multiple actions, such as picking up or delivering bicycles. After completion of a task, the driver requests a new one. Therefore, for each vehicle, a route must be created by assigning the next task to be performed, unaware of which tasks will follow. We begin by presenting the dynamic framework, which outlines the steps involved in creating the next route, starting with the driver's request for the next task and ending with the execution of that task. Thereafter, we explain briefly the real-time station status, the demand predictions, and the no-service penalties.

4.1. Dynamic framework

The steps of creating the next route are presented in the dynamic framework (Figure 4.1). The process begins in the upper-left corner, where a driver requests a new task after completing the previous one. This request enters the platform, which analyzes the real-time status of all stations defined as the number of bicycles available and free docking spots. Additionally, demand predictions are generated for the upcoming period. The system status and demand predictions are used as inputs to the no-service level function. This value is computed for each station using the function $f(S, D)$, where S represents the current inventory level and D denotes the predicted demand. Together, these values represent the expected inventory at each station if no-repositioning actions are performed. The resulting no-service penalties are then used as inputs for the Dynamic Bicycle Repositioning Problem with Flexible Time Horizon (DBRPFTH). This model determines which stations to visit, how many bicycles to relocate, and the most efficient route to follow. A detailed description of the model is provided in Chapter 5. Once the next route is determined, the task is assigned to the driver for execution. Upon completion, the driver submits a new task request, and the cycle repeats.

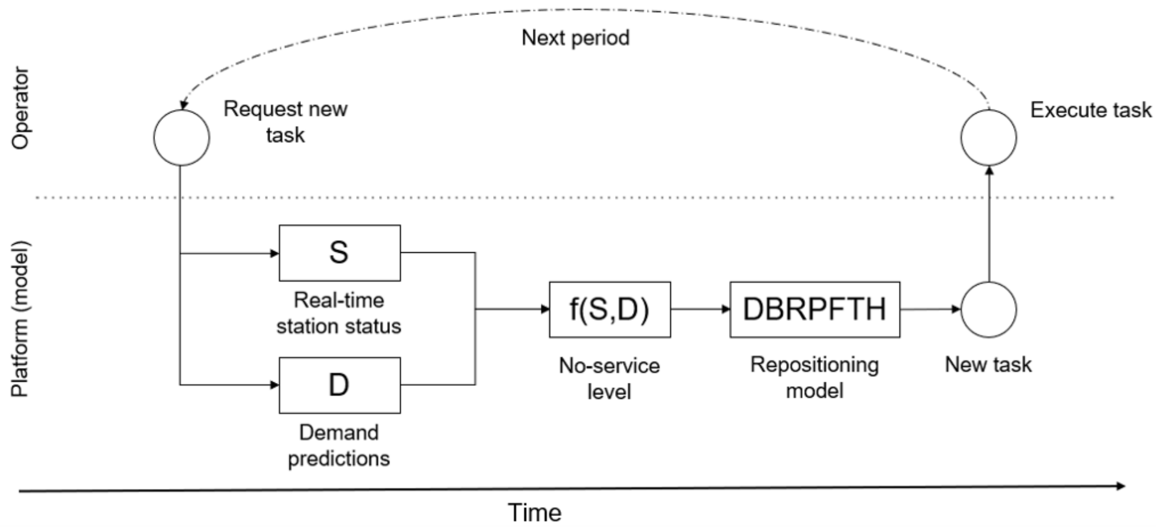

Figure 4.1: Schematic overview of the dynamic framework

4.2. Real-time station status

The real-time station status describes the current state at each station. It contains the number of available bicycles and the number of free docking spots. The data will be collected from open-access data sources in the standard General Bikeshare Feed Specification (GBFS) format, which contains real-time information on shared bicycle systems. This includes data on vehicle and station locations, dock availability, bicycle characteristics, service pricing, and rental conditions. It is updated every 30 seconds.

4.3. Demand predictions

The demand predictions estimate fluctuations in bicycle levels at each station for the next time interval. Hence, the expected inventory level strongly depends on the demand predictions. To ensure flexibility, the prediction horizon adapts to the planning of each repositioning route. During peak hours, shorter forecasts of 30 minutes are preferred due to rapid changes in station inventory levels. In contrast, longer prediction windows of 40 to 45 minutes are more suitable for overnight and off-peak periods when fluctuations are less frequent. Additionally, predictions align dynamically with route schedules. For example, if a route ends at 10:40 with half-hour forecast horizon, predictions are generated for 10:40–11:10 instead of using a fixed 10:30–11:00 window. This is achieved by partitioning predictions into five-minute intervals, allowing flexible adaptation to different start times and planning horizons.

The prediction model relies solely on unconstrained historical demand data, which means that changes in bicycle inventory are unaffected by station capacity limits. If a station has no bicycles and is still empty in the next period, the actual demand might be nonzero, but any demand is unobservable. Similarly, when a station reaches full capacity, the number of users who wanted to return a bicycle, but were unable to do so remains unknown. To ensure accuracy, we only used historical data where the station inventory remained within capacity limits during the next interval.

The demand prediction model consists of several steps. The first step involves clustering past dates based on their demand patterns, without considering calendar characteristics. The number of trips for all stations will be aggregated to an overall demand for a day, so it will be clustered to a total demand. Note that if the demand pattern for several stations differs significantly from the usual pattern, the day will be in a different cluster. Importantly, rainy days are excluded from clustering because rain negatively affects demand, potentially leading to misclassifications.

In the second step, a machine learning (ML) classification model predicts the cluster to which a future date belongs based on its calendar characteristics, such as the day of the week, month, holidays, and special events.

In the third step, an ML regression model forecasts demand for the next time interval using multiple inputs:

- The assigned cluster for the date
- Lagged variables from previous periods
- Demand fluctuations at similar past dates within the same cluster
- Calendar variables
- The latest meteorological forecasts, including rain, temperature, wind speed, and humidity

To improve reliability, three separate models are trained: one to predict the mean demand, and two quantile regressors to estimate the lower and upper demand bounds. These quantile models are trained with loss functions that emphasize penalties on either negative demand (more bicycles taken than returned) or positive demand (more bicycles returned than taken), which allow an estimation of the 20th and 80th percentiles, respectively. To improve robustness against demand uncertainty, the system adopts a worst-case scenario approach based on the estimated demand bounds. When the lower and upper quantiles indicate that two additional bicycles might be picked up or dropped off and the station is currently empty, the scenario with more pickups is considered the worst case. In contrast, if the station is full, the scenario with more deliveries is considered worse. This conservative approach helps mitigate the risk of demand fluctuations, ensuring a more reliable and stable system.

4.4. No-service penalties

The no-service penalties will be determined using real-time station status and demand predictions, as these two factors allow us to determine the expected inventory level for each station in the upcoming period, assuming no repositioning actions are taken. The ideal scenario for each station is a balance between the number of bicycles and free docks, ensuring that the station is neither empty nor full, and is capable of accommodating future demand. Since future demand is uncertain, stations that are near-empty or near-full are considered suboptimal. A no-service penalty will be imposed if the station occupancy becomes nearly empty or nearly full. Operators often apply priority rules, where certain stations, such as those in city centers, near key buildings, or at major transit hubs, are given higher priority. To reflect this, priority stations will have a higher penalty factor.

The problem definition outlined in this chapter forms the basis for the mathematical formulation presented in the next chapter.

5
Model formulation

In this chapter, we introduce the Dynamic Bicycle Repositioning Problem with Flexible Time Horizon (DBRPFTH), which focuses on determining the stations to visit, the optimal number of bicycles to relocate, and creating an efficient route with a flexible time horizon. We formulate the DBRPFTH as a standard Mixed Integer Linear Program (MILP) to optimize user satisfaction, while ensuring operational efficiency. This model is based on the Static Bicycle Repositioning Problem (SBRP) introduced by Raviv et al. (2013). We begin by introducing the measures for operational efficiency and user satisfaction before we delve deeper into the model.

5.1. Measures

In bicycle repositioning problems, the goal is typically to find a balance between operational efficiency, often expressed in terms of duration, distance, or CO_2 emissions, and user satisfaction, which is commonly measured by the level of service. In this thesis, we introduce specific measures for both operational efficiency and user satisfaction in the context of a continuously operating bicycle system with changing conditions throughout the day.

5.1.1. Operational efficiency

Operational efficiency is important in evaluating repositioning strategies with the primary goal of creating efficient routes that respond effectively to fluctuations in station inventory levels. In our context, where the number of drivers is predetermined and vehicles operate continuously, traditional cost-minimization objectives, such as reducing driver wages and vehicle operations, are less relevant. Since drivers request new routes immediately after completing the previous one, minimizing routing costs does not significantly impact efficiency. Instead, we focus on creating efficient routes to maximize user satisfaction.

We introduce a maximum route duration T, which determines how long a vehicle can operate before requesting a new route. The route duration is influenced by the travel time between stations, which varies throughout the day, and the stop time, which includes both a fixed time to pick up or deliver bicycles and an additional component based on the number of bicycles being moved. To account for variations in travel time and system conditions throughout the day, we divide the day into three distinct time intervals: peak hours, off-peak hours, and overnight. Each time period falls within one of these intervals, allowing parameters to be defined at the interval level rather than for each individual time period. To further adapt to these variations, we use different duration matrices for travel times in each interval, ensuring that travel times accurately reflect the conditions during repositioning routes.

In order to address variations in demand and station inventory fluctuations throughout the day, we introduce the time-dependent maximum route duration, T_t, where t represents the specific time period. During peak hours, when demand is high and station inventory fluctuates significantly, it becomes crucial to respond quickly to these fluctuations by planning shorter routes. To achieve this, we reduce the maximum route duration T_t during peak hours, ensuring that repositioning routes can adapt faster to rapid changes in inventory levels. On the other hand, during overnight and off-peak hours, when demand is lower and station inventory levels are more stable, the model shifts its focus towards optimizing operational efficiency. In these periods, longer routes can be planned because the station inventory fluctuates less, allowing more actions to be efficiently included in a single route. To accommodate this, the maximum route duration T_t is extended, which allows the system to optimize for fewer, but longer, repositioning tasks.

5.1.2. User satisfaction

User satisfaction is assessed by no-service penalties, as described in Section 4.4. A penalty is imposed if the expected station occupancy, determined by the current inventory level and the demand forecast for the upcoming period, falls below a lower threshold (nearly empty) or exceeds an upper threshold (nearly full). These penalties are calculated based on the squared deviation from the thresholds, which places greater emphasis on larger deviations. This reflects the fact that extreme imbalances can result in unmet demand. Moreover, the impact of repositioning a single bicycle can go beyond the benefit of a single user, as it may enable a chain of successful rentals and returns (Chemla et al., 2013). Additionally, a fixed penalty is applied if a station becomes completely empty or full. Together, these mechanisms discourage extreme inventory levels and promote more balanced stations, ultimately aiming to improve service levels and reduce the risk of long-term negative consequences.

To determine the optimal inventory level for each station in the upcoming periods, demand forecasts are incorporated for both the upcoming period and the period after. The demand in the upcoming period reflects the demand during which repositioning takes place, while the forecast for the two periods ahead represents the expected demand after repositioning has occurred. For example, if a station is expected to become (nearly) full in the coming period, we know that a pickup is needed, but it remains unclear whether this should involve removing 25%, 50%, or 75% of the bicycles. The forecast for the period after repositioning helps refine this decision: If demand is expected to be positive (more returns than withdrawals), it is preferable to leave fewer bicycles in the station; if demand is expected to be negative (more withdrawals than returns), having more bicycles available is advantageous. This component is integrated into a quadratic penalty function, where the penalty is lowest at the target inventory level that best aligns with anticipated demand. A relatively small weight is assigned to this term to ensure that inventory levels outside the acceptable range are still penalized more heavily than those within it. This structure prioritizes meeting immediate demand predictions while using future forecasts to fine-tune inventory within the preferred bounds.

Lastly, higher weights are applied to the penalties of priority stations to ensure that they receive more attention in optimization.

5.2. Notation

The dynamic repositioning problem is described by the following set and parameters:

N	Set of nodes, including the depot, indexed by $i = 0, \dots, \lvert N \rvert$.
i_o	Origin node, where the repositioning operation starts.
s_i^0	Number of bicycles at node i before the repositioning operation starts.
d_i^1	Predicted demand at node i for the next period.
d_i^2	Predicted demand at node i for two periods ahead.
c_i	Number of docks installed at node i, referred to as the node's capacity.
c^{max}	Maximum number of docks.
l	Initial load in the vehicle.
k	Capacity (number of bicycles) of the vehicle.
$f_i(s_i)$	Penalty function reflecting user satisfaction for node i.

r_{ij}	Traveling time from node i to node j.
T_t	Repositioning time, i.e., maximum duration of a repositioning route in period t.
V	Constant time required for visiting a node.
P	Time required to pick up a bicycle from a node.
D	Time required to deliver a bicycle to a node.
M	Upper bound on number of arcs, default is $\|N\|$

Note that only imbalanced stations are considered to reduce computational complexity (Regue & Recker, 2014). The depot is assumed to have no demand and capacity, because the vehicle cannot drop-off or pick up bicycles at the depot. Since the driver requests a new route after completing the previous one, the vehicle does not necessarily start its route from the depot. Instead, the starting location is referred to as the origin node, denoted by i_o. Station capacity is defined as the total number of usable bicycles and docks, excluding any that are broken or unavailable.

The following decision variables will be used:

x_{ij}	Binary variable which equals one if the vehicle travels directly from node i to node j, and zero otherwise.
y_{ij}	Number of bicycles carried on the vehicle when it travels directly from node i to node j. If the vehicle does not travel directly from i to j, y_{ij} is zero.
y_i^P	Number of bicycles picked up at node i.
y_i^D	Number of bicycles delivered at node i.
z_i^P	Binary variable which equals one if bicycles are picked up from node i, and zero otherwise.
z_i^D	Binary variable which equals one if bicycles are delivered to node i, and zero otherwise.
z_i	Binary variable which equals one if action is performed at node i, and zero otherwise.
q_i	Auxiliary variable used for sub-tour elimination constraints.
s_i	Expected inventory level at node i at the end of the repositioning operation.

5.3. Objective function

Our objective is to maximize user satisfaction by minimizing no-service penalties, using the satisfaction measures defined in Subsection 5.1.2. This approach is supported by Caggiani and Ottomanelli (2013), who model utility as a plateau function, where utility increases from an empty station to a sufficient lower bound, remains high within an optimal range, and decreases as the station nears full capacity. However, in contrast to their approach, we aim to determine the optimal inventory level within the acceptable range (the plateau) based on the demand predictions for the period after repositioning, which corresponds to two periods ahead.

We introduce the penalty function to evaluate the expected inventory levels. This function combines a convex component and a piecewise-linear component to reflect the penalties related to imbalanced stations, based on the research of Raviv et al. (2013). To incorporate this function into an optimization model, we reformulate both components into linear terms using auxiliary variables and constraints. This transformation enables the complete model to be expressed as a Mixed-Integer Linear Program (MILP), following the approach of Raviv et al. (2013). It is important to note that the values of $f_i(s_i)$ can be calculated from the input variables.

The function contains the following parameters:

L	Percentage of station capacity considered as almost empty
H	Percentage of station capacity considered as almost full
p^L	Penalty (almost) empty station
p^H	Penalty (almost) full station
p^{cL}	Constant penalty empty station
p^{cH}	Constant penalty full station
α_i	Weight/scaling factor for priority stations, set to 1 if station i is not a priority
β	Weight/scaling factor for optimal inventory level

The penalty function for station $i \in N$ given station inventory level s_i is as follows:

$$\begin{aligned} f_i(s_i) = \alpha_i \Bigg(& p^L \big(\max(c_i \cdot L - s_i, 0)\big)^2 + p^H \big(\max(s_i - c_i \cdot H, 0)\big)^2 \\ & + \beta \left(s_i - \left(\frac{L+H}{2} \cdot c_i - d_i^2 \right) \right)^2 \\ & + p^{cL} \, 1_{s_i \le 0} + p^{cH} \, 1_{s_i \ge c_i} \Bigg) \quad \forall i \in N \end{aligned} \tag{5.1}$$

In the repositioning model, scalar factors α_i are introduced to assign greater importance to penalties at priority stations. The penalty function consists of five components that can be divided into two types: convex components and piecewise linear components. The first three components form the convex part of the function, while the last two components introduce piecewise linear penalties. Consequently, the penalty function $f_i(s_i)$ can be separated into a convex penalty function $f_i^C(s_i)$ and a piecewise linear penalty function $f_i^L(s_i)$.

$$\begin{aligned} f_i^C(s_i) = \alpha_i \Bigg(& p^L \big(\max(c_i \cdot L - s_i, 0)\big)^2 + p^H \big(\max(s_i - c_i \cdot H, 0)\big)^2 \\ & + \beta \left(s_i - \left(\frac{L+H}{2} \cdot c_i - d_i^2 \right) \right)^2 \Bigg) \quad \forall i \in N \\ f_i^L(s_i) = \alpha_i \Big(& p^{cL} \, 1_{s_i \le 0} + p^{cH} \, 1_{s_i \ge c_i} \Big) \quad \forall i \in N \end{aligned}$$

The first component of the convex penalty function addresses situations where the inventory level of a station s_i drops below a low threshold, indicating that the station is nearly empty. This penalty is squared and weighted by the scalar p^L, increasing the penalty as the inventory level is further from the threshold. The second component focuses on stations that are (nearly) full, applying a similar squared penalty with a scalar p^H as s_i approaches the station's full capacity.

The third component is always active, but is specifically used in determining the optimal inventory level within the acceptable range. It distinguishes between inventory levels that are still balanced, between the lower and upper bounds, by using the demand forecast for two periods ahead, which is the demand in the period after repositioning. If the post-repositioning forecast is positive (indicating more returns than withdrawals), it is better to leave fewer bicycles at the station to prevent overflow. As a result, the penalty function then favors lower inventory levels. In contrast, if the forecast is negative (more withdrawals than returns), the system benefits from higher bicycle availability, and the penalty is minimized at higher inventory levels. The minimum penalty corresponds to the inventory level that best matches the expected post-repositioning demand. A small scalar weight β is applied to this term to ensure that inventory levels outside the boundaries are always worse than inventory levels within the boundaries. This ensures that the model prioritizes meeting immediate demand (during repositioning). Future forecasts are used to fine-tune inventory levels within the bounds. This results in more balanced and forward-looking repositioning decisions.

The piecewise linear components introduce additional penalties when a station becomes either completely empty or completely full. These components are captured by an indicator function that equals 1 if s_i is non-positive (station is empty) or greater than or equal to the station's capacity c_i (station is full). A negative inventory level is possible when the station expects a large demand deficit, i.e. significantly more bicycles are taken than returned. Similarly, an inventory level exceeding the station's capacity can occur when the station expects a large demand surplus, i.e. significantly more bicycles are returned than taken.

Together, the above mentioned five components ensure that the repositioning strategy not only addresses immediate demand and balances stations within acceptable inventory levels, but also accounts for future demand fluctuations, minimizes the likelihood of stations becoming entirely empty or full, and prioritizes critical stations that directly impact service levels.

To transform the convex penalty function f_i^C into linear terms, we must first determine the domain of the expected inventory level s_i. This level depends on the initial inventory before repositioning s_i^0, the predicted demand d_i^1, and the repositioning actions y_i^P (pickups) and y_i^D (drop-offs). The initial inventory s_i^0 ranges from 0 to the station capacity c_i, while d_i^1 can take any integer value. Repositioning actions y_i^P and y_i^D are limited by the number of bicycles at the station or the free docks available before repositioning. If d_i^1 is zero, the inventory s_i will remain within the range $[0, c_i]$. However, if d_i^1 is negative, which means that more bicycles are taken than returned, s_i can drop below zero, extending the lower bound of the domain to d_i^1. Conversely, if d_i^1 is positive, which means that more bicycles are returned than taken, s_i can exceed the station's capacity, extending the upper bound to $d_i^1 + c_i$. Thus, the overall domain of s_i is given by:

$$[min(0, d_i^1), max(c_i, c_i + d_i^1)]$$

To replace the convex penalty functions f_i^C with a linear term and linear constraints, we introduce the following set and equations:

$$\begin{aligned} U &= [min(0, d_i^1), max(c_i - 1, c_i - 1 + d_i^1)] & \\ b_{iu} &\equiv f_i^C(u+1) - f_i^C(u) & \forall i \in N, u \in U \\ a_{iu} &\equiv f_i^C(u) - b_{iu} \cdot u & \forall i \in N, u \in U \end{aligned}$$

The set U represents the indices over which s_i is defined. The term b_{iu} captures the marginal penalty associated with adding the $(u+1)^{th}$ bicycle to station i. Together, a_{iu} and b_{iu} represent the intercept and slope, respectively, of the linear function that approximates the convex penalty function f_i^C at the u^{th} level. These linear approximations are included as constraints in the MILP to ensure that the correct value of the convex part is considered in the objective function, depending on the expected inventory level s_i. The resulting value of the convex penalty function for each station is captured by the variable g_i, which is included in the objective function.

To linearize the piecewise penalty function f_i^L, we replace the indicator functions in $f^L(\cdot)$ with binary decision variables. Specifically, we define the binary variables w_i and o_i, which indicate whether station i is empty or full, respectively. That is, $w_i = 1$ if the inventory level $s_i \leq 0$, and $o_i = 1$ if the inventory level $s_i \geq c_i$. To enforce this behavior, we introduce two linear constraints using a sufficiently large constant C.

5.4. Mathematical model

The mathematical model for the dynamic bicycle repositioning problem is as follows:

$$
\begin{aligned}
\min \quad & \sum_{i\in N} g_i + \alpha_i(p^{cL}\cdot w_i + p^{cH}\cdot o_i) && && (5.2)\\
\text{s.t.} \quad & g_i \geq a_{iu} + b_{iu}\cdot s_i && \forall i\in N, u\in U && (5.3)\\
& s_i > -C\cdot w_i && \forall i\in N && (5.4)\\
& s_i < c_i + C\cdot o_i && \forall i\in N && (5.5)\\
& s_i = s_i^0 + d_i^1 - (y_i^P - y_i^D) && \forall i\in N && (5.6)\\
& y_i^P - y_i^D = \sum_{j\in N, j\neq i} y_{ij} - \sum_{j\in N, j\neq i} y_{ji} && \forall i\in N\setminus i_o && (5.7)\\
& l + y_{i_o}^P - y_{i_o}^D = \sum_{j\in N, j\neq i_o} y_{i_o j} && && (5.8)\\
& y_{ij} \leq k\cdot x_{ij} && \forall i,j\in N, i\neq j && (5.9)\\
& \sum_{j\in N, j\neq i} x_{ij} = \sum_{j\in N, j\neq i} x_{ji} && \forall i\in N && (5.10)\\
& y_i^P \leq s_i^0 && \forall i\in N && (5.11)\\
& y_i^D \leq c_i - s_i^0 && \forall i\in N && (5.12)\\
& \sum_{i\in N}\Big(\sum_{j\in N} r_{ij}\cdot x_{ij} + P\cdot y_i^P + D\cdot y_i^D + V\cdot z_i\Big) \leq T_t && && (5.13)\\
& q_j \geq q_i + 1 - M(1 - x_{ij}) && \forall i\in N, j\in N\setminus i_o, i\neq j && (5.14)\\
& \frac{y_i^P}{c^{max}} \leq z_i^P \leq y_i^P && \forall i\in N && (5.15)\\
& \frac{y_i^D}{c^{max}} \leq z_i^D \leq y_i^D && \forall i\in N && (5.16)\\
& z_i^P + z_i^D <= 1 && \forall i\in N && (5.17)\\
& \frac{z_i^P + z_i^D}{2} \leq z_i \leq z_i^P + z_i^D && \forall i\in N && (5.18)\\
& x_{ij} \in \mathbb{B} && \forall i,j\in N && (5.19)\\
& z_i^P, z_i^D, z_i \in \mathbb{B} && \forall i\in N && (5.20)\\
& y_i^P, y_i^D \in \mathbb{Z}_{\geq 0} && \forall i\in N && (5.21)\\
& y_{ij} \geq 0 && \forall i,j\in N, i\neq j && (5.22)\\
& q_i \geq 0 && \forall i\in N && (5.23)
\end{aligned}
$$

The objective function (5.2) minimizes the no-service penalties, as described by (5.1). Constraints (5.3) determine the convex no-service penalties based on the linearization intercept and slope, while constraints (5.4) and (5.5) set the piecewise-linear no-service penalties. The inventory balance is enforced by constraints (5.6), which set the expected inventory levels after repositioning. Additionally, bicycle flow is conserved by the constraints (5.7), while the constraint (5.8) ensures that the flow leaving the origin node equals the initial vehicle load plus any actions taken at the origin. The constraints (5.9) limit the number of bicycles carried by a vehicle to its capacity and ensure that no bicycles are transported along an arc if the vehicle does not use that arc.

The vehicle flow is preserved by the constraints (5.10), while the number of bicycles picked up or delivered to a station is restricted by the constraints (5.11) and (5.12), ensuring that the pickups do not exceed the available inventory and the deliveries do not exceed the remaining dock capacity. The constraints (5.11) inherently enforce the non-negativity of the inventory variables, while the constraints (5.12) ensure that the inventory at each station and the depot remains within capacity limits. As a result, explicit constraints for inventory non-negativity and capacity limits are not required.

The maximum route duration is constrained by (5.13). Note that the return to the origin node is set to zero to enforce a round-trip structure in the model, while the actual route will not return. The sub-tour elimination is enforced through constraints (5.14), which prevent small loops in the route. Constraints (5.15) and (5.16) ensure that an action is only performed if bicycles are picked up or delivered. The constraints (5.17) enforce that vehicles pick up or deliver bicycles at a station, but not both. Furthermore, constraints (5.18) ensure that stations are only visited when an action is performed.

Binary and integer constraints on decision variables are imposed by (5.19)-(5.21), while non-negativity conditions are enforced by constraints (5.22) and (5.23). Non-negativity constraints on s_i cannot be enforced, as the expected inventory level can be negative after repositioning when the demand deficit (d_i^1) is highly negative. Similarly, s_i may exceed the station capacity c_i when demand is exceptionally high. In such cases, it is assumed that unmet demand is lost, but it is still taken into account in the objective function to quantify its impact. Finally, the integrality of y_{ij} and s_i is implicitly ensured by the integer nature of y_i^P, y_i^D and s_i^0.

Similar to Raviv et al. (2013) we introduce valid inequalities to reduce computation time.

$$\sum_{j \in N} x_{i_o j} \geq 1 \tag{5.24}$$

$$\sum_{j \in N, j \neq i} x_{ij} \leq 1 \qquad \forall i \in N \tag{5.25}$$

$$y_i^P \leq min(s_i^0, k) \sum_{j \in N} x_{ij} \qquad \forall i \in N \tag{5.26}$$

$$y_i^D \leq min(c_i - s_i^0, k) \sum_{j \in N} x_{ij} \qquad \forall i \in N \tag{5.27}$$

$$y_i^P + y_i^D \geq \sum_{j \in N} x_{ij} \qquad \forall i \in N \setminus \{i_o, 0\} \tag{5.28}$$

The constraint (5.24) ensures that the vehicle leaves the origin at least once. In the LP relaxation, this departure count can take fractional values. By enforcing the departure of a complete vehicle, the vehicle flow conservation constraints (5.10) guarantee that this condition holds for all visited nodes. The constraints (5.25) ensure that each station is visited at most once. In addition, constraints (5.26) and (5.27) impose a stricter form of constraints (5.11) and (5.12), which limit the number of bicycles picked up or delivered based on vehicle capacity and apply only when the vehicle visits the station. Finally, constraints (5.28) ensure that vehicles visit a station only if bicycles are picked up or delivered.

6

Model extensions

In this chapter, we present multiple extensions to the model we previously introduced. First, we expand the model to optimize a system with a mixed bicycle fleet. In addition, we explore the inclusion of vehicle locations in the model. These extensions introduce both complexity and flexibility, enabling more precise optimization of bicycle repositioning strategies.

6.1. Mixed bicycle fleet

As electric bicycles become more prevalent and mixed bicycle-sharing systems emerge, the model is extended to accommodate both mechanical and electric bicycles. The integration of electric bicycles into bicycle sharing systems enhances accessibility, reduces physical effort, and enables longer trips, making cycling a more attractive option (Ghamami & Shojaei, 2018). Electric bicycles also improve the competitiveness of bicycle sharing systems with private cars, as they offer greater speed and convenience, particularly on hilly routes (Bieliński et al., 2021). User preferences further highlight the benefits of electric bicycles, as they are more resilient to long distances, high temperatures, and poor air quality, although precipitation remains a limiting factor (Campbell et al., 2016). Given these advantages, hybridization of mechanical and electric bicycles is expected to become a global trend, allowing operators to balance affordability with enhanced service quality (Shui & Szeto, 2020). Existing dynamic bicycle relocation problems (DBRPs) primarily address multi-vehicle scenarios but do not consider the complexities introduced by a mixed bicycle fleet (Shui & Szeto, 2020). When dealing with DBRPs that incorporate a mixed bicycle fleet, additional operational constraints must be considered. Specifically, loading and unloading strategies become more complex, as an excess of one bicycle type may not effectively resolve shortages of another (Shui & Szeto, 2020).

We present an extended model designed to accommodate a system with both mechanical and electric bicycles. We assume that all docks are available for both types of bicycles, that each bicycle type occupies the same amount of space in the repositioning vehicle, and that the time required to pick up or deliver a mechanical or electric bicycle is identical. To incorporate a mixed bicycle fleet, we extend the model by introducing new decision variables. The pickup and delivery variables, originally denoted as y_i^P and y_i^D, are now replaced by $y_i^{P_m}$, $y_i^{P_e}$, $y_i^{D_m}$ and $y_i^{D_e}$, where m refers to mechanical bicycles and e to electric bicycles. Similarly, the routing variable y_{ij} is expanded to y_{ij}^m and y_{ij}^e. The station inventory variable s_i is now represented separately as s_i^m and s_i^e, while the predicted demands d_i^1 and d_i^2 are split into $d_i^{1_m}$, $d_i^{1_e}$, $d_i^{2_m}$ and $d_i^{2_e}$. Additionally, the initial inventory level s_i^0 is replaced by $s_i^{0_m}$ and $s_i^{0_e}$. The initial vehicle load, previously denoted as l, is now differentiated as l_m and l_e. Finally, the pickup and delivery indicators z_i^P and z_i^D are expanded into $z_i^{P_m}$, $z_i^{P_e}$, $z_i^{D_m}$ and $z_i^{D_e}$, ensuring that the model accounts for both bicycle types separately.

The convex penalty function $f_i^C(\cdot)$ consists of five components: a penalty for almost no mechanical bicycles, a penalty for almost no electric bicycles, a penalty for almost full station, and two quadratic penalties for future demand forecasts. Furthermore, the parameter L is replaced by L_m and L_e.

$$f_i^{L_m}(s_i^m) = \alpha_i(p^L(max(c_i \cdot L_m - s_i^m, 0))^2) \qquad \forall i \in N \tag{6.1}$$
$$f_i^{L_e}(s_i^e) = \alpha_i(p^L(max(c_i \cdot L_e - s_i^e, 0))^2) \qquad \forall i \in N \tag{6.2}$$
$$f_i^{H}(s_i^m + s_i^e) = \alpha_i(p^H(max((s_i^m + s_i^e) - c_i \cdot H, 0))^2) \qquad \forall i \in N \tag{6.3}$$
$$f_i^{IB_m}(s_i^m) = \alpha_i(\beta \cdot (s_i^m - (\frac{L_m + 0.5 \cdot H}{2} \cdot c_i - d_i^{2_m}))^2) \qquad \forall i \in N \tag{6.4}$$
$$f_i^{IB_e}(s_i^e) = \alpha_i(\beta \cdot (s_i^e - (\frac{L_e + 0.5 \cdot H}{2} \cdot c_i - d_i^{2_e}))^2) \qquad \forall i \in N \tag{6.5}$$

The convex penalty components can be linearized similar to the original model. The domains for $f_i^{L_m}$, $f_i^{L_e}$ and f_i^H vary because the functions have different inputs. To address these differences, we define the sets U_1, U_2 and U_3 as the respective domains for $f_i^{L_m}$ and $f_i^{IB_m}$, $f_i^{L_e}$ and $f_i^{IB_e}(s_i^e)$, and f_i^H. Note that the optimal inventory level has been modified such that the optimal inventory level for s_i^m and s_i^e together is equal to the original optimal inventory level.

$$U_1 = [min(0, d_i^{1_m}), max(c_i - 1, c_i - 1 + d_i^{1_m})]$$
$$U_2 = [min(0, d_i^{1_e}), max(c_i - 1, c_i - 1 + d_i^{1_e}]$$
$$U_3 = [min(0, d_i^{1_m} + d_i^{1_e}), max(c_i - 1, c_i - 1 + d_i^{1_m} + d_i^{1_e})]$$

With the sets defined above, the constraints (5.3) will be replaced by:

$$g_i^{L_m} \geq a_{iu_1}^{L_m} + b_{iu_1}^{L_m} \cdot s_i^m \qquad \forall i \in N, u_1 \in U_1 \tag{6.6}$$
$$g_i^{L_e} \geq a_{iu_2}^{L_e} + b_{iu_2}^{L_e} \cdot s_i^e \qquad \forall i \in N, u_2 \in U_2 \tag{6.7}$$
$$g_i^{H} \geq a_{iu_3}^{H} + b_{iu_3}^{H} \cdot (s_i^m + s_i^e) \qquad \forall i \in N, u_3 \in U_3 \tag{6.8}$$
$$g_i^{IB_m} \geq a_{iu_1}^{IB_m} + b_{iu_1}^{IB_m} \cdot s_i^m \qquad \forall i \in N, u_1 \in U_1 \tag{6.9}$$
$$g_i^{IB_e} \geq a_{iu_2}^{IB_e} + b_{iu_2}^{IB_e} \cdot s_i^e \qquad \forall i \in N, u_2 \in U_2 \tag{6.10}$$

The piecewise line penalty function $f_i^L(\cdot)$ consists of three components: a penalty for the absence of mechanical bicycles, a penalty for the absence of electric bicycles and a penalty for full stations.

$$f_i^L(s_i^m, s_i^e) = \alpha_i \left(p^{cL}\, 1_{s_i^m \leq 0} + p^{cL}\, 1_{s_i^e \leq 0} + p^{cH}\, 1_{s_i^m + s_i^e \geq c_i} \right) \qquad \forall i \in N \tag{6.11}$$

These components can be linearized in a similar way to those in the original formulation. The decision variables w_i are replaced by w_i^m and w_i^e respectively. As a result, the original constraints (5.4) and (5.5) are modified and replaced by the following updated constraints:

$$s_i^m > -C \cdot w_i^m \qquad \forall i \in N \tag{6.12}$$
$$s_i^e > -C \cdot w_i^e \qquad \forall i \in N \tag{6.13}$$
$$s_i^m + s_i^e < c_i + C \cdot o_i \qquad \forall i \in N \tag{6.14}$$

The objective function (5.2), which now incorporates the modified decision variables, must be updated accordingly. In addition, all constraints involving these variables or related parameters require adjustments. This includes constraints (5.6)-(5.9), (5.11)-(5.13), and (5.15)-(5.18), as well as (5.20)-(5.22).

$$\min \quad \sum_{i \in N} g_i^{L_m} + g_i^{L_e} + g_i^{H} + g_i^{IB_m} + g_i^{IB_e} + \alpha_i(p^{cL} \cdot (w_i^m + w_i^e) + p^{cH} \cdot o_i) \tag{6.15}$$

$$s_i^m = s_i^{0_m} + d_i^{1_m} - (y_i^{P_m} - y_i^{D_m}) \quad \forall i \in N \quad (6.16)$$
$$s_i^e = s_i^{0_e} + d_i^{1_e} - (y_i^{P_e} - y_i^{D_e}) \quad \forall i \in N \quad (6.17)$$
$$y_i^{P_m} - y_i^{D_m} = \sum_{j \in N, j \neq i} y_{ij}^m - \sum_{j \in N, j \neq i} y_{ji}^m \quad \forall i \in N \setminus i_o \quad (6.18)$$
$$y_i^{P_e} - y_i^{D_e} = \sum_{j \in N, j \neq i} y_{ij}^e - \sum_{j \in N, j \neq i} y_{ji}^e \quad \forall i \in N \setminus i_o \quad (6.19)$$
$$l_m + y_{i_o}^{P_m} - y_{i_o}^{D_m} = \sum_{j \in N, j \neq i_o} y_{i_o j}^m \quad (6.20)$$
$$l_e + y_{i_o}^{P_e} - y_{i_o}^{D_e} = \sum_{j \in N, j \neq i_o} y_{i_o j}^e \quad (6.21)$$
$$y_{ij}^m + y_{ij}^e \leq k \cdot x_{ij} \quad \forall i, j \in N, i \neq j \quad (6.22)$$
$$y_i^{P_m} \leq s_i^{0_m} \quad \forall i \in N \quad (6.23)$$
$$y_i^{P_e} \leq s_i^{0_e} \quad \forall i \in N \quad (6.24)$$
$$y_i^{D_m} + y_i^{D_e} \leq c_i - (s_i^{0_m} + s_i^{0_e}) \quad \forall i \in N \quad (6.25)$$
$$\sum_{i \in N} (\sum_{j \in N} r_{ij} \cdot x_{ij} + P \cdot (y_i^{P_m} + y_i^{P_e}) + D \cdot (y_i^{D_m} + y_i^{D_e}) + V \cdot z_i) \leq T_t \quad (6.26)$$
$$\frac{y_i^{P_m}}{c^{max}} \leq z_i^{P_m} \leq y_i^{P_m} \quad \forall i \in N \quad (6.27)$$
$$\frac{y_i^{P_e}}{c^{max}} \leq z_i^{P_e} \leq y_i^{P_e} \quad \forall i \in N \quad (6.28)$$
$$\frac{y_i^{D_m}}{c^{max}} \leq z_i^{D_m} \leq y_i^{D_m} \quad \forall i \in N \quad (6.29)$$
$$\frac{y_i^{D_e}}{c^{max}} \leq z_i^{D_e} \leq y_i^{D_e} \quad \forall i \in N \quad (6.30)$$
$$z_i^{P_m} + z_i^{D_m} <= 1 \quad \forall i \in N \quad (6.31)$$
$$z_i^{P_e} + z_i^{D_e} <= 1 \quad \forall i \in N \quad (6.32)$$
$$\frac{z_i^{P_m} + z_i^{P_e} + z_i^{D_m} + z_i^{D_e}}{4} \leq z_i \leq z_i^{P_m} + z_i^{P_e} + z_i^{D_m} + z_i^{D_e} \quad \forall i \in N \quad (6.33)$$
$$z_i^{P_m}, z_i^{P_e}, z_i^{D_m}, z_i^{D_e}, z_i \in \mathbb{B} \quad \forall i \in N \quad (6.34)$$
$$y_i^{P_m}, y_i^{P_e}, y_i^{D_m}, y_i^{D_e} \in \mathbb{Z}_{\geq 0} \quad \forall i \in N \quad (6.35)$$
$$y_{ij}^m, y_{ij}^e \geq 0 \quad \forall i, j \in N, i \neq j \quad (6.36)$$

Furthermore, the valid inequalities (5.26)-(5.28) can be updated.

$$y_i^{P_m} \leq min(s_i^{0_m}, k) \sum_{j \in N} x_{ij} \quad \forall i \in N \quad (6.37)$$
$$y_i^{P_e} \leq min(s_i^{0_e}, k) \sum_{j \in N} x_{ij} \quad \forall i \in N \quad (6.38)$$
$$y_i^{D_m} + y_i^{D_e} \leq min(c_i - (s_i^{0_m} + s_i^{0_e}), k) \sum_{j \in N} x_{ij} \quad \forall i \in N \quad (6.39)$$
$$y_i^{P_m} + y_i^{P_e} + y_i^{D_m} + y_i^{D_e} \geq \sum_{j \in N} x_{ij} \quad \forall i \in N \setminus \{i_o, 0\} \quad (6.40)$$

Note that in constraints (6.37) and (6.38), x_{ij} is 1 as well if bicycles of the other type are picked up.

6.2. Vehicle-aware repositioning

To enhance operational efficiency and reduce computational complexity, various strategies have been proposed. A widely used method involves clustering stations into fixed service areas (Jiménez-Meroño & Soriguera, 2024). For example, O'Mahony and Shmoys (2015) apply clustering during peak hours to direct vehicles to stations with similar behavior, while Wang et al. (2022) propose a demand-driven clustering approach. Although clustering helps reduce problem size, it limits vehicle flexibility by imposing fixed spatial boundaries. As an alternative, vehicle-aware repositioning incorporates the expected end locations of all repositioning vehicles directly into the decision-making process. This look-ahead approach enables the model to assign tasks in a way that anticipates future opportunities (Jiménez-Meroño & Soriguera, 2024; Pfrommer et al., 2014).

In our approach, when a driver requests a new task, we explicitly consider the ongoing tasks of currently busy vehicles. If a busy vehicle is expected to finish its current task in a location close to a particular demand, that location is not assigned to the requesting driver. We limit potential task assignments by excluding a fixed number of locations near the expected end locations of busy vehicles. This vehicle-aware coordination improves fleet efficiency and scalability by flexibly limiting certain locations, without setting fixed area boundaries. Importantly, no such restrictions are applied to locations near the current position of the vehicle requesting the task.

Figure 6.1 shows how vehicle-aware repositioning affects the set of stations considered for assignment in Zaragoza. In this example, four repositioning vehicles operate simultaneously. The green vehicle marks the position of the driver currently requesting a task. The three red vehicles represent the expected end locations of currently busy vehicles. Stations shown as green dots are the closest locations to the requesting vehicle, while red dots indicate stations located near the expected end points of busy vehicles. These stations are temporarily restricted for the requesting vehicle. Blue dots represent all other stations that remain available. The visual representation shows how certain areas of the network are effectively "reserved" for other vehicles that are better positioned to serve them in the near future.

In the scenario with ten restricted locations, some overlap occurs between green and red dots, meaning certain stations are both close to the requesting vehicle and within a restricted area. Proximity is calculated using travel time rather than straight-line (Euclidean) distance, so stations that appear geographically close on the map may still be farther away in practice. The model also excludes stations already assigned to other vehicles, stations previously visited by the requesting vehicle, and stations that are currently balanced. Consequently, the number of restricted locations directly influences the set of stations considered in the optimization.

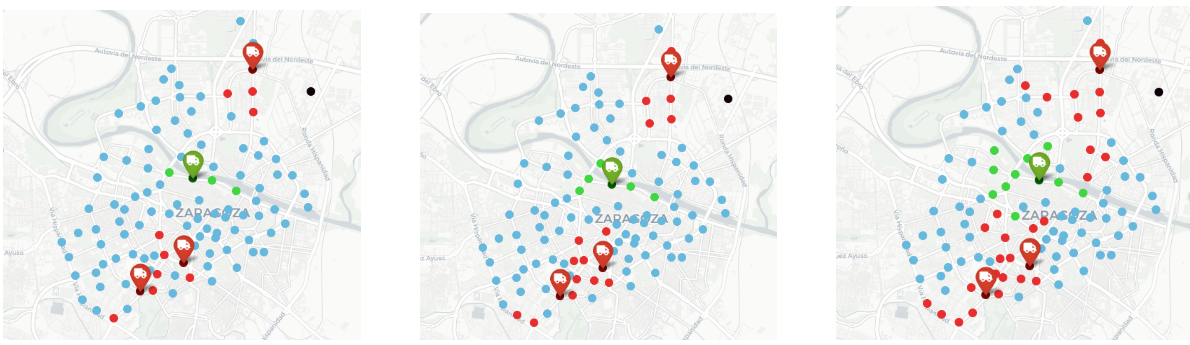

(a) 3 restricted locations **(b)** 5 restricted locations **(c)** 10 restricted locations

Figure 6.1: Vehicle-aware repositioning

• = Depot • = Current vehicle • = Stations near current vehicle • = End locations of busy vehicles
• = Stations near end locations of busy vehicles • = Other stations

7 Simulation

To evaluate the performance of bicycle repositioning strategies, we will perform a series of simulation experiments comparing four operational approaches: (1) no repositioning, (2) a static repositioning strategy, (3) the proposed state-of-the-art dynamic repositioning strategy developed in this study and (4) a perfect information approach. In the static repositioning strategy, bicycles are redistributed only during the night shift, when demand is relatively low. The purpose of this strategy is to prepare the system for the following day. This can be seen as a simplified variant of the proposed dynamic repositioning approach, restricted to nighttime operations with a larger number of repositioning vehicles. As a result, the expected demand during the day is not explicitly considered, only the demand that occurs during the repositioning period itself is taken into account.

The perfect information approach serves as a benchmark by assuming that the dynamic repositioning strategy is executed with full knowledge of future demand. Instead of relying on forecasts, the actual demand for the upcoming planning horizon is provided to the strategy. This allows us to determine the maximum achievable satisfied demand under the proposed dynamic approach if the forecasts were perfectly accurate. Any remaining unsatisfied demand can then be attributed to resource limitations, such as the number or capacity of repositioning vehicles, available bicycles, or dock space, rather than to the repositioning strategy itself. This highlights the performance limits of the dynamic strategy and clarifies the extent to which lost demand is caused by factors other than forecast accuracy.

7.1. Demand

To replicate real-world system dynamics, user demand is simulated on the basis of historical station status data. Origin–destination (OD) trips are generated by Poisson sampling, which models discrete trip events and ensures system stability. This approach avoids the imbalances that can occur when modeling arrivals and departures separately or when applying the Skellam distribution to represent net demand, since any positive net demand at one station must be exactly offset by a negative value at another station.

The expected demand for each OD pair in a given hour is calculated as the average demand from historical data. Similar to demand predictions, periods with constrained data must be excluded from historical data. Such constrained periods occur when the stations are empty or full, making the actual demand unobservable during these times. For each OD pair-hour combination, we count the five-minute periods during which the departure or arrival station is constrained. A departure station is only considered constrained when it is empty (as a full station does not restrict departures), while an arrival station is constrained only when it is full. Demand for these constrained periods is determined based on the level of available unconstrained data. If more than half of an hour consists of unconstrained periods, the constrained demand is estimated proportionally, assuming that demand would have been linear relative to the observed unconstrained periods. However, if more than half of the hour is constrained, the remaining unconstrained data is considered unreliable. In such cases, the demand is replaced by the average hourly demand for that specific OD pair.

Since many OD pair combinations are rare events, either because stops are located very close to each other or too far apart, there are many OD pairs for which no demand was observed in the historical data. As a result, the OD matrix of the expected demand contains many zeros. This would imply that the Poisson-distributed demand for these OD pairs is always zero, preventing the simulation of any trips along these pairs. However, such rare trips do occasionally occur and therefore a small probability should be assigned. To address this issue, Laplace smoothing is applied. This technique assigns a small positive value to all OD pairs to prevent zero probabilities. To ensure that the total expected demand remains unchanged, the probabilities are subsequently rescaled. Specifically, for each OD pair (i, j), the smoothed probability is calculated as:

$$P_{ij}^{smoothed} = \frac{\lambda_{ij} + \gamma}{\lambda_{total} + \gamma V}$$

where:

- λ_{ij} is the original expected demand for OD pair (i, j)
- λ_{total} is the total expected demand over all OD pairs
- V is the total number of OD pairs
- γ is the smoothing parameter

The smoothed expected demand for OD pair (i, j) is then scaled back to ensure that the total demand remains unchanged:

$$\hat{\lambda}_{ij} = \lambda_{total} \cdot P_{ij}^{smoothed} = \lambda_{total} \cdot \frac{\lambda_{ij} + \gamma}{\lambda_{total} + \gamma V}$$

This approach guarantees that the sum of all smoothed demands remains equal to the original total demand, while allowing all OD pairs to remain possible in simulations.

7.2. Experiment

The system dynamics of the shared bicycle service will be simulated over a complete day to capture realistic variations in user demand throughout different time periods, including peak hours, off-peak periods, and overnight times. Each simulation day will begin at night to minimize the influence of the initial system state. We perform 10 simulation runs for each scenario to obtain stable results. The simulation incorporates three repositioning shifts per day, each lasting eight hours: a night shift (22:00–06:00), a morning shift (06:00–14:00), and an afternoon shift (14:00–22:00). To ensure smoother operations, the start and end times of some vehicles within a shift may be slightly adjusted. For example, the start of a shift for part of the fleet may be delayed by an hour to avoid all vehicles being at the depot simultaneously. Each vehicle starts its shift by leaving the depot and returns there at the end of the shift. During operations, drivers request a new task after completing their previous one. The next task can only include stations that are unbalanced, not visited by other vehicles or visited in the previous route of the vehicle. The repositioning algorithm then determines the next assignment. If no feasible route is found at that moment, the driver will wait for five minutes before requesting a new task. Note that the route duration can be lower than the maximum route duration, consequently the driver will request a new task sooner. The system is updated every minute, taking into account both user trips and repositioning movements. When a bicycle trip and a repositioning action occur at the same station within the same minute, the repositioning trip is prioritized. If no bicycles of the requested type are available at the departure station, the user cannot start the trip, and the request is recorded as unsatisfied demand. An important assumption is that users do not switch bicycle types if their preferred type is unavailable. This implies that a shortage of one bicycle type cannot be compensated by a surplus of another. If no available docks are found at the intended destination station, the user parks the bicycle at the nearest station with available docks. However, this is still counted as unsatisfied demand at the intended destination station. Note that when the simulation day begins, there may already be bicycles moving, trips that started before the simulation's start time but arrive after it has begun. To account for these trips, historical trip data is used to represent bicycles that were actually moving at the time the system status is retrieved. These are not simulated movements but real trips that started before the simulation period and are scheduled to arrive after its start, ensuring the initial system status reflects the true state of bicycles on the network. When a vehicle arrives at a station to pick up or deliver bicycles, but the required number of bicycles (for pickup) or free docks (for delivery) is not available, the vehicle adjusts by picking up or delivering as many as possible. This limitation may affect the feasibility of subsequent planned actions, as the number of bicycles on the vehicle may differ from what was initially intended. In such cases, the algorithm first checks whether the remaining bicycle stock on the vehicle allows the original plan to continue. If not, the planned quantities for the following stations are adjusted accordingly or the remaining actions are canceled.

The standard analysis will focus on a typical weekday under normal conditions. Both single-type and mixed bicycle fleets will be evaluated to assess how different fleet types influence repositioning performance. Additionally, the simulation will investigate the effects of several factors, including day types (weekdays, weekends, and holidays), restrictions on vehicle end locations, initial system states, fleet size, vehicle capacity, and varying levels of demand levels. Each simulation scenario will be assessed using the validation metrics introduced in Section 9.1, with particular emphasis on user satisfaction and operational efficiency. This comprehensive approach will provide insights in the performance of different repositioning strategies under various operating conditions.

8
Case Study

We evaluate the performance of our models through simulation experiments using data from two station-based bicycle-sharing systems in Zaragoza and Valladolid. The data is collected from open-access data sources with standard General Bikeshare Feed Specification (GBFS) formatted data, which contain real-time information on shared bicycle operations. This includes details on vehicle and station locations, dock availability, bicycle types, service pricing, and rental conditions. The predictions and simulated demand used in our models are generated as described in Section 4.3 and Chapter 7, based on historical GBFS data from both systems.

8.1. System characteristics

Table 8.1 provides an overview of the key characteristics of the bicycle-sharing systems in Zaragoza and Valladolid for daily operation. The most notable difference between the two systems is the number of users, with Zaragoza having approximately seven times more users than Valladolid. Additionally, Zaragoza's system exclusively operates electric bicycles, while Valladolid's system includes both mechanical and electric bicycles. Although the number of stations is similar in both cities, there is a significant difference in the total number of docks. However, the difference in docks per station between the two systems is quite small. The total bicycle fleet is comparable across both systems, with Valladolid's fleet being nearly evenly divided between mechanical and electric bicycles, whereas Zaragoza's fleet consists entirely of electric bicycles.

The daily trips-to-bicycle ratio is considerably higher in Zaragoza, indicating that bicycles are used more frequently. Moreover, in Valladolid, electric bicycles are preferred to mechanical ones. Both systems have similar docks-to-bicycles ratio, which reflects the likelihood that users find an available dock when returning a bicycle. According to the literature, an ideal docks-to-bicycles ratio for efficient bicycle repositioning is approximately 2 (Soriguera & Jiménez-Meroño, 2020). For example, Brescia has a ratio of 1.9 (Angelelli et al., 2022), while Barcelona has a ratio of 2.7. Higher ratios are often recommended in less optimized systems (Soriguera & Jiménez-Meroño, 2020).

Table 8.1: System characteristics

	Zaragoza	Valladolid
Number of daily users	13430	1815
Number of bicycle types	1	2
Number of stations	108	100
Number of docks	2137	1846
Docks/Station ratio	19.8	18.5
Number of mechanical bicycles	-	420
Number of electric bicycles	960	415
Total number of bicycles	960	835
Trips/bicycle ratio	14.0	2.2
Docks/Bicycles ratio	2.2	2.2

8.2. System parameters

Table 8.2 summarizes the main parameter settings used in the simulation. While some settings are shared across both systems, others are system-specific. Both Zaragoza and Valladolid operate continuously throughout the day, with demand varying over time periods. During the night (22:00–06:00) the demand is minimal, while during peak periods (07:00–09:00, 13:00–15:00, 17:00–19:00) there is high demand. The maximum route duration is adjusted accordingly: 45 minutes during the night, 40 minutes during off-peak hours, and 30 minutes during peak periods. The first task of each shift has a maximum route duration of 45 minutes to allow the vehicle to depart from the depot. If the remaining shift time is less than the maximum route duration plus the time needed to return to the depot, the vehicle must return to the depot to ensure it finishes its shift on time. A fixed stop-time of 5 minutes is applied for each repositioning task, with an additional 0.5 minutes per bicycle moved. The penalty parameters are mostly consistent across both systems. The scalar p^L for (near) empty stations is higher than p^H for (near) full stations, as empty stations are considered more problematic. Denied departures often mean users leave the system, while in the case of full stations, users continue their trip by diverging to nearby stations. The inventory bounds between both systems differ slightly: Zaragoza uses fixed bounds of 0.2 (lower) and 0.8 (upper), while Valladolid applies bounds of 0.15 and 0.7, defined per vehicle type. This results in a higher combined lower bound but a lower upper bound in Valladolid, providing more flexibility to adjust station inventories in line with its lower system activity. The smoothing parameter γ used for demand estimation is set higher for Zaragoza to reflect its larger volume of trips.

Table 8.2 also shows the differences in vehicle operations. In Zaragoza two vehicles are used at night, four during the morning and three during the afternoon. In Valladolid one vehicle is deployed at night and four during the day. In the static approach, four vehicles are used in both systems during the night to prepare the system for the upcoming day. Vehicle capacities differ as well: Vehicles in Zaragoza carry up to 22 bicycles, while in Valladolid the capacity is between 11 and 14.

Table 8.2: System parameter settings

Category	Parameter	Zaragoza	Valladolid
Number of vehicles	Number of vehicles (night)	2	1
	Number of vehicles (morning)	4	4
	Number of vehicles (afternoon)	3	4
Vehicle restrictions	Vehicle capacity	22	11-14
	Restricted locations	0	0
Duration (min)	Maximum duration (night)	45	45
	Maximum duration (off-peak)	40	40
	Maximum duration (peak)	30	30
Stop-time (min)	Constant stop-time	5	5
	Stop-time per bicycle	0.5	0.5
Penalties	α	5	5
	β	0.001	0.001
	Lower bound	0.2	0.15
	Upper bound	0.8	0.7
	p^L	1	1
	p^H	0.5	0.5
	p^{cL}	20	20
	p^{cH}	20	20
Simulation	γ	0.001	0.0002

9

Results

Results are obtained by applying the simulation methods explained in Chapter 7 to the different models explained in Chapter 5 and Chapter 6 to the case studies introduced in Chapter 8. All optimization problems are solved using Gurobi 12.0.1, with an average computation time of 5 seconds per task for both systems. For each scenario, 10 simulation runs are performed to obtain stable results, and the results are presented as the average values along with their standard deviations. First, we provide an overview of the validation measures, followed by an analysis of the simulation results for the different repositioning strategies applied to both the single-type system (Zaragoza) and the mixed-fleet system (Valladolid). Finally, we analyze the effects of different factors on the results.

9.1. Validation

The main Key Performance Indicator (KPI) for assessing the effectiveness of a bicycle repositioning strategy is the satisfied demand during the day, which reflects the number of users who can find an available bicycle at their origin and a free dock at their destination. This metric directly captures user satisfaction, which serves as the main objective in repositioning shared bicycles. Similarly, the amount of unsatisfied demand can be determined, both at the origin (empty station) and at the destination (full station), possibly indicating a shortage of bicycles or docks. Note that no-service penalty is not a good KPI to compare different repositioning strategies, since the no-repositioning approach has no routes, and the static repositioning strategy repositioning only takes place during the night, when the penalties are generally lower.

In addition to satisfied demand, several operational KPIs are essential for evaluating the efficiency of the repositioning process. They can be grouped into two main categories. The first category relates to station performance and includes: the number of empty and full stations, and the number of bicycles present at stations. The second category concerns repositioning operations and includes: the number of tasks, the number of stations visited, the number of bicycles repositioned, and the average route duration (measured in seconds).

The definition of service level can vary depending on how a bicycle sharing system is managed. Some systems focus on immediate availability. These systems aim to prevent stations from becoming empty or full in real time or minimize the time a station is empty or full (Brinkmann et al., 2015). These systems react to real-time conditions with minimal use of demand predictions. Others take a more forward-looking approach, planning repositioning based on anticipated demand to ensure a consistent user experience over time. Because of these different strategies, the same situation, such as a temporarily empty station, might be seen as a big problem in one system but acceptable in another if more bicycles are expected to arrive soon. Although the relocation model remains the same, the definition of service level must be adapted to align with the specific operational goals of each system, either focusing on current imbalances or future inventory.

9.2. Single-type system

The average results for the single-type system in Zaragoza are presented in Table 9.1. As expected, the scenario without repositioning results in the lowest level of satisfied demand, although the system is still able to meet more than 80% of the total demand. While implementing a static repositioning strategy results in a modest improvement in demand satisfaction, the dynamic strategy leads to a substantial increase, serving around 430 additional users over the day. With perfect information of future demand, the system's performance could improve even further, potentially satisfying 300 more users (2.4%) beyond what the regular dynamic strategy achieves.

Across all strategies, unsatisfied demand is considerably higher at trip origins than at destinations. Even the perfect information scenario results in almost 10% of trips being denied due to a lack of available bicycles. This suggests that the main limitation of the system is the unavailability of bicycles at the origin stations, rather than the lack of docking space at destinations. This observation is supported by the higher average number of empty stations compared to full stations throughout the day. Even with perfect information, an average of 7.5 stations remain empty, highlighting the challenge of maintaining sufficient bicycle availability throughout the network. Interestingly, although the dynamic approach results in more full stations than the static approach, it leads to lower unsatisfied demand at destinations. This suggests that by repositioning bicycles based on demand predictions, more trips can be fulfilled despite the higher number of full stations. The number of bicycles at stations reflects their overall availability. During operation, bicycles are constantly moving, either because users are making trips or because drivers are repositioning them. In the dynamic approach (with perfect information), the availability of bicycles at stations is 20–30 bicycles lower than in the no-repositioning scenario. This means that, on average, 20-30 more bicycles are moving in the system. This increase is partly due to more users finding a bicycle at their trip origin and partly due to repositioning actions.

In terms of repositioning effort, the static strategy results in a relatively high number of stations visited per task compared to the dynamic strategy. A possible explanation for this is the longer maximum route duration during the night, which allows for longer repositioning routes. Meanwhile, the dynamic strategy moves more bicycles per task on average, as fluctuations in station inventory levels during the day require larger repositioning movements. Interestingly, the total number of tasks, stations visited, and bicycles moved when deploying the regular dynamic strategy is very similar to those in the dynamic repositioning with perfect information. This highlights that, without additional operational effort, improvements in demand prediction accuracy enable higher demand satisfaction through better decisions. Finally, the average route duration for the static strategy is substantially higher, since all repositioning occurs during nighttime hours when longer routes are allowed. In contrast, dynamic and perfect information strategies yield similar and shorter average route durations, indicating that route length is not a distinguishing factor between these two approaches.

Table 9.1: Average results for single-type system

Type	No repositioning	Static	Dynamic	Perfect information
Demand satisfaction				
Satisfied demand	11086.9 (105.1)	11283.3 (115.3)	11691.8 (110.9)	12008.9 (112)
Unsatisfied demand origin	1691.4 (106.3)	1624 (86.4)	1351.4 (71.2)	1162.7 (93.8)
Unsatisfied demand destination	696.4 (57.6)	567.3 (55.6)	428.2 (50.9)	295.7 (52)
Satisfaction rate (%)	82.3	83.7	86.8	89.1
Station status				
Empty stations	12.8 (0.4)	9.5 (0.3)	8.8 (0.4)	7.5 (0.4)
Full stations	9.1 (0.5)	4.4 (0.4)	4.8 (0.8)	3.9 (0.6)
Number of bicycles at stations	862 (0.9)	842 (3.3)	840.4 (2.1)	835.3 (4.1)
Repositioning				
Number of tasks	0	43.8 (0.4)	114.3 (1.2)	113.8 (1.3)
Number of stations visited	0	132.9 (3.6)	278.7 (4)	274.4 (5.2)
Number of bicycles repositioned	0	233.8 (12.3)	698.2 (24.8)	701 (30.4)
Route duration (s)	-	2599 (6)	2248.5 (24.7)	2257.9 (21.2)

In order to analyze the repositioning strategies in more detail, we first look at the demand pattern and the number of bicycles at stations throughout the day (Figure 9.1). Since these patterns are generally similar across all strategies and the differences are subtle, we show them only for the dynamic approach. Demand is almost zero overnight and starts to increase around 6:00 AM. The morning peak occurs between 7:00 AM and 9:00 AM, followed by a sharp drop until about 1:00 PM. A second, larger peak appears between 2:00 PM and 3:00 PM, after which demand stays high until around 8:00 PM before falling again. The highest levels of unsatisfied demand occur during these peak hours. Throughout the day, unsatisfied demand remains consistently higher at origin stations compared to destination stations. Furthermore, we observe that the confidence interval, represented by the shaded area around the mean, is narrow, indicating a high consistency and low variability in the results. The number of bicycles at stations follows the opposite trend of demand, fewer bicycles are at stations when demand is high, as more trips are being made. This effect is amplified during the day when more repositioning vehicles are active, meaning more bicycles are repositioned. This partly explains why unsatisfied demand at origin stations is particularly high during peak hours.

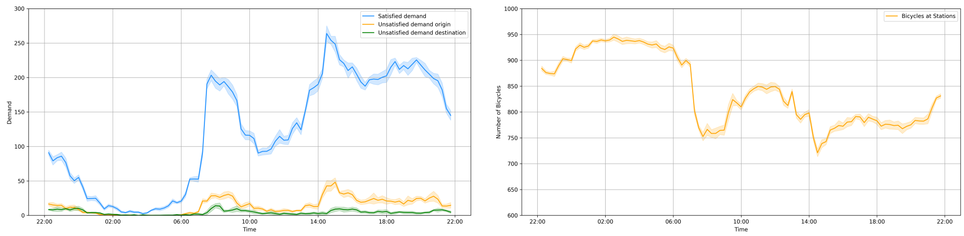

(a) Demand **(b)** Number of Bicycles at Stations

Figure 9.1: Analysis over the Day

Figure 9.2 shows station status over the day for each repositioning strategy. A clear difference is visible at night: without repositioning, the number of empty and full stations stays high, since minimal trip activity means the system state does not change. We also observe that there are more full stations than empty ones at the start of the night, and that the static approach reduces the number of critical stations more quickly, which is expected given the use of four vehicles. During peak hours, however, the no-repositioning and static strategies perform similarly to the dynamic approach in terms of empty stations. This is likely because there are already few bicycles available, and demand fluctuations during these hours are too large for the repositioning vehicles to handle, causing major imbalances in a short time. After the morning peak, especially between 11:00 AM and 2:00 PM, the benefits of the dynamic approach become clearer, as it helps reduce the number of critical stations more effectively. Finally, even with perfect information, the reduction in empty and full stations is limited, highlighting the difficulty of fully balancing the system under real-world conditions.

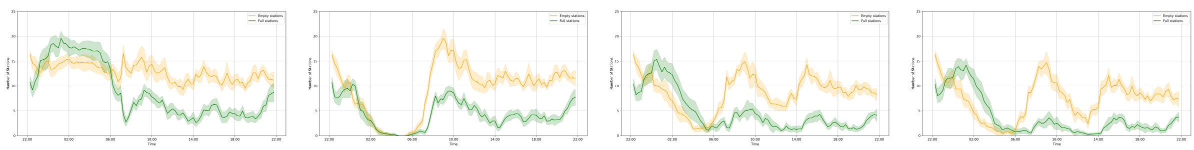

(a) No repositioning **(b)** Static repositioning **(c)** Dynamic repositioning **(d)** Perfect information

Figure 9.2: Station Status over the Day

9.3. Mixed-fleet system

The average results for mixed-fleet system in Valladolid are presented in Table 9.2. As expected, the scenario without repositioning results in the lowest level of satisfied demand, with the system serving around 1540 trips. The static repositioning strategy already improves this by approximately 80 additional trips, while the dynamic strategy achieves a substantial further increase, satisfying 120 more users and covering over 95% of total demand. With perfect information about future demand, the system could only serve about 40 additional trips (2%) compared to the regular dynamic strategy, suggesting that the dynamic approach already comes close to the maximum achievable performance.

Similar to the single-type system, the unsatisfied demand at the trip origins is notably higher than at the destinations for all strategies. This again indicates that the primary reason for unsatisfied demand is the lack of bicycles available rather than the lack of dock space. Notably, for the dynamic strategy and perfect information scenarios the unsatisfied demand at the destinations is reduced to almost zero. The distinction between mechanical and electric bicycles further reveals that stations experience electric empty states more frequently than mechanical empty states, highlighting the greater challenge of maintaining the availability of electric bicycles. With the dynamic approach, an average of 3.9 stations are empty with regard to electric bicycles compared to only 0.7 stations lacking mechanical bicycles. This disparity aligns with the considerably higher demand for electric bicycles, which is more than 4 times higher than the demand for mechanical bicycles. The number of mechanical and electric bicycles at stations is similar across the different strategies, which is expected given that the number of trip is much lower than in single-type system in Zaragoza. However, the difference between strategies is slightly larger for electric bicycles (around 10) than for mechanical bicycles (around 5). This is consistent with the higher usage of electric bicycles, which probably also increases the need for their repositioning.

With regard to repositioning movements, the static strategy leads to more visited stations per task compared to the dynamic strategy, probably because the allowed route duration is longer during the night period. The number of bicycles moved per station is relatively low for all strategies, indicating many small-scale repositioning moves. The total number of tasks, stations visited, and bicycles moved in the regular dynamic approach and the dynamic approach with perfect information is very similar, implying that further improvements in demand prediction could increase satisfied demand without additional operational effort. Finally, the average route duration for the static strategy is considerably higher than for the dynamic and perfect information strategies, as expected, since it takes place only during the night. The dynamic and perfect information strategies yield nearly identical route durations, indicating that route length is not a differentiating factor between these approaches.

Table 9.2: Average results for mixed-fleet system

Type	No repositioning	Static	Dynamic	Perfect information
Demand satisfaction				
Satisfied demand	1542.6 (61.3)	1622.8 (53.5)	1743.5 (47.7)	1782.7 (45.1)
Unsatisfied demand origin	250.2 (21.8)	178 (18.6)	79.6 (9.1)	42.7 (9.8)
Unsatisfied demand destination	36.2 (16)	27.5 (11.2)	3.3 (2.3)	0.6 (0.8)
Satisfaction rate (%)	84.3	88.8	95.5	97.6
Station status				
Empty stations mechanical bicycles	5.6 (0.8)	2.6 (0.8)	0.7 (0.2)	0.3 (0.2)
Empty stations electric bicycles	13.1 (1.2)	9 (0.7)	3.9 (0.5)	2.4 (0.7)
Full stations	1.8 (0.3)	0.7 (0.3)	0.1 (0.1)	0 (0)
Mechanical bicycles at stations	415.1 (0.2)	412.5 (1.3)	411.6 (0.9)	410.8 (1.2)
Electric bicycles at stations	400.2 (0.4)	396.8 (1)	391.1 (3)	391.5 (3)
Repositioning				
Number of tasks	0	44 (0)	115.1 (1)	116 (1.1)
Number of stations visited	0	136.6 (2.3)	296.4 (3.8)	299.8 (4.2)
Number of bicycles repositioned	0	180.9 (8.8)	451.4 (17.2)	474.3 (10)
Route duration (s)	-	2570.2 (17.8)	2197.2 (24.1)	2180 (19.1)

To analyze the repositioning strategies in more detail, we first look at the demand pattern and the number of bicycles at stations throughout the day (Figure 9.3). As in the single-type system, these patterns are very similar across all strategies, with only subtle differences. Therefore, we present them only for the dynamic approach. Demand stays relatively high until midnight, after which it drops sharply to nearly zero. Around 6:00 AM, demand begins to rise again, reaching a morning peak between 7:00 and 9:00 AM, then declining significantly until about 1:00 PM. The highest demand peak occurs between 2:00 and 3:00 PM, after which demand remains elevated until midnight. The results show that almost all unsatisfied demand is avoided, and repositioning helps stabilize the system following peak periods. Notably, unsatisfied demand at destination stations is barely visible. Similar to the single-type case, the narrow confidence interval indicates high consistency and low variability in the outcomes.

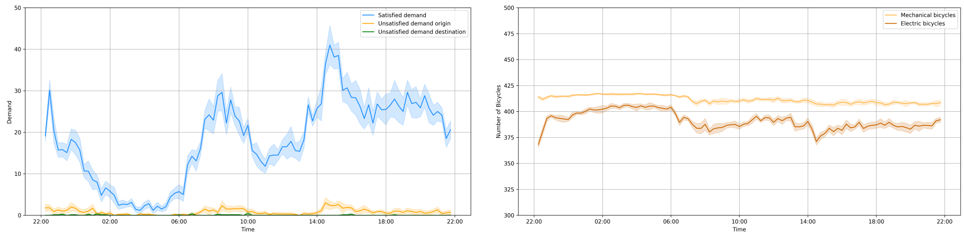

(a) Demand **(b)** Number of Bicycles at Stations

Figure 9.3: Analysis over the Day

Figure 9.4 shows how station status changes over the day for each repositioning strategy. There is a significant difference between the no-repositioning and static strategies on one side, and the dynamic approach (with perfect information) on the other. In the no-repositioning and static cases, where no repositioning is carried out during the day, the number of empty stations steadily rises, with particularly many stations lacking electric bicycles. Full stations are also present throughout the day. In contrast, the dynamic strategies keep the number of empty stations low during the day, with further reductions after peak periods and overnight. Full stations are almost negligible under these approaches. Although perfect information allows for slightly better handling after peak hours, its advantage over regular dynamic repositioning is relatively small.

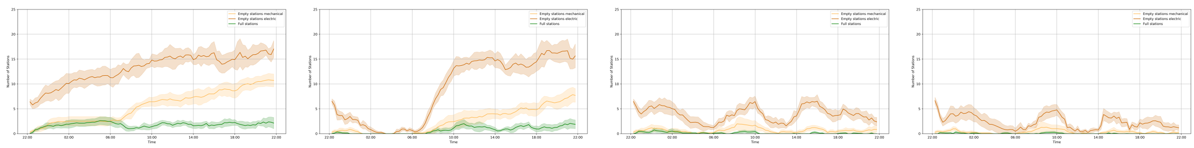

(a) No repositioning **(b)** Static repositioning **(c)** Dynamic repositioning **(d)** Perfect information

Figure 9.4: Station Status over the Day

9.4. Sensitivity analysis of system parameters

To gain a deeper understanding of how various model parameters affect system performance, an additional sensitivity analysis was performed on the dynamic repositioning strategy. First, the initial system status was found to have a negligible effect on satisfied demand, indicating that the main findings are robust to variations in starting conditions. The results for a single-type system are presented in Figure 9.5. The light blue bars represent the absolute number of satisfied demand, while the dark blue line indicates the satisfaction rate. It is worth noting that the absolute satisfied demand may increase even when the satisfaction rate decreases, and vice versa.

First of all, we observe that adjusting the number of vehicles by one, either by removing or adding, leads to a change of about 1% in satisfaction rate, with a decrease when removing a vehicle and an increase when adding one. Changing vehicle capacity has a negligible effect on satisfied demand. The impact of station prioritization was also evaluated. In practice, 13 priority stations are designated by the operator and assigned higher penalties when unbalanced. Removing these priorities offers more flexibility to balance the network, but the effect on the satisfaction rate is negligible. Incorporating demand predictions has a small impact, only slightly bringing the results closer to those obtained under perfect information. Satisfaction rates remain similar with vehicle-aware repositioning regardless of the number of restricted locations. However, the computation time decreases from 2.7 to 1.3 seconds, and this effect is likely to be more significant in larger systems. Weekend operations were analyzed, during which the fleet size is reduced to two vehicles during the day and one at night. Weekend and festive days have similar demand patterns that are significantly lower than weekdays. This lower demand makes repositioning easier, which results in a higher percentage of satisfied demand. Finally, as the system demand increases, the absolute number of satisfied trips increases significantly, while the satisfaction rate drops only by 1% to 4%, remaining above 80%. This suggests that the system can accommodate substantial demand growth with limited loss in service quality.

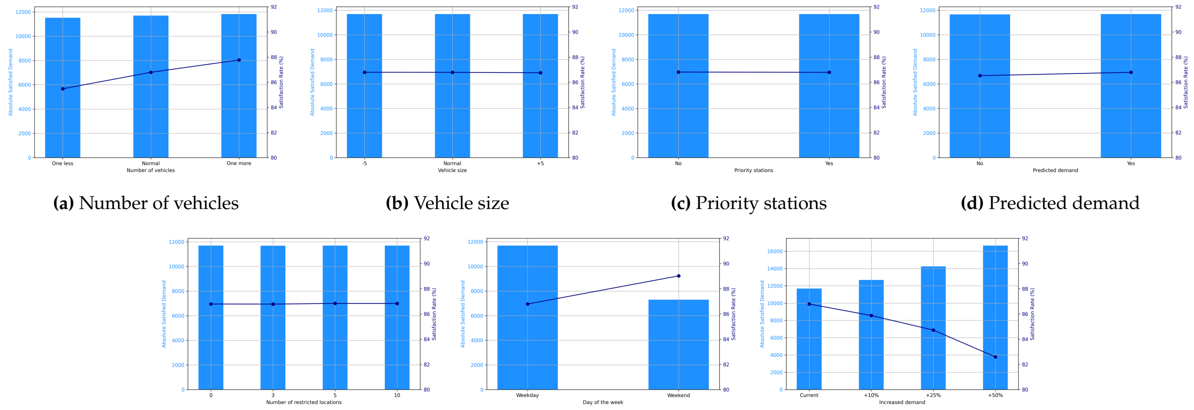

(a) Number of vehicles **(b)** Vehicle size **(c)** Priority stations **(d)** Predicted demand

(e) Vehicle-aware repositioning **(f)** Day of the week **(g)** Increased demand

Figure 9.5: Sensitivity Analysis of a Single-Type System

In Valladolid, there are no priority stations. Therefore, prioritization is not further investigated. First, we find that the initial system status has minimal impact on satisfied demand, confirming the robustness of the results to variations in starting conditions, consistent with the findings for the single-type fleet. Figure 9.6 presents the sensitivity analysis for the mixed-fleet system. Increasing the number of vehicles improves the satisfaction rate from 94% to 96%, with the largest gain observed when increasing from fewer vehicles to the current number. Varying vehicle size has a small effect on performance; even when decreasing vehicle capacity to just 9 bicycles, the satisfaction rate remains considerably high. The small differences observed are mainly due to randomness of demand. Incorporating demand predictions improves model performance, increasing the satisfaction rate by almost 1% and bringing it within 2% of the perfect information benchmark. Vehicle-aware repositioning lowers satisfaction slightly but reduces the computation time from 5 to 3 seconds. Its lower performance is likely due to the relatively small number of unbalanced stations, meaning that restricting the set of candidate stations greatly limits the solution space. As in Zaragoza's single-type system, weekend demand is much lower than during the week. However, unlike Zaragoza, satisfaction rates drop during weekends, suggesting that the number of repositioning vehicles may be insufficient to maintain similar service levels. Finally, when the system demand increases, the number of trips increases substantially, while the satisfaction rate decreases only moderately (by 1%–4%), remaining above 90%. This indicates that the system can handle significant demand growth with minimal loss in service quality.

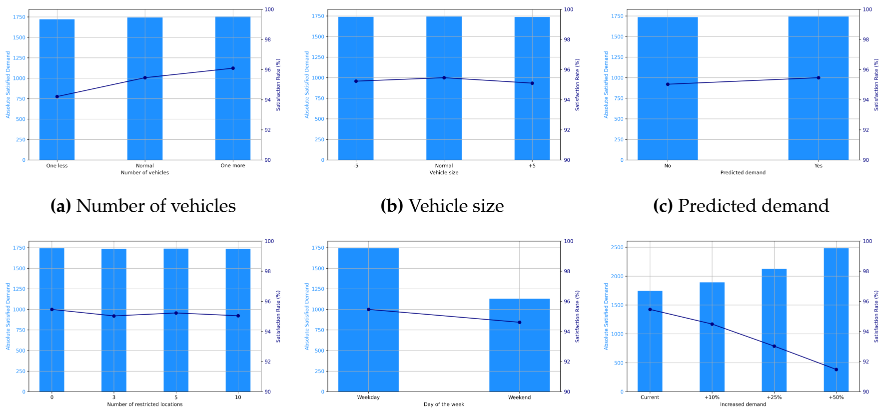

(a) Number of vehicles **(b)** Vehicle size **(c)** Predicted demand

(d) Vehicle-aware repositioning **(e)** Day of the week **(f)** Increased demand

Figure 9.6: Sensitivity Analysis of a Mixed-Fleet System

9.5. Analysis of empty stations

In the results, we observed significantly more unsatisfied demand at trip origins, more empty stations than full ones, and that even with perfect information it was difficult to avoid empty stations during the day. To investigate the root of this issue, we consider two options: increasing the penalties for empty stations to prioritize them, or adding bicycles to test whether a shortage of bicycles is the underlying cause. Figure 9.7 shows the analysis of empty stations in the single-type system. Adjusting the scalar for nearly empty stations, increasing the constant for empty stations, or applying both changes does not improve satisfaction rates. In contrast, adding 200 extra bicycles increases satisfaction by 1%. However adding more bicycles does not lead to further improvement, and adding 300 or 400 bicycles actually reduces satisfaction because it increase the number of full stations. The optimal level is an addition of 200 bicycles, which leads to an improvement of 1%. Nevertheless, the number of critical stations (empty + full) decreases only marginally, suggesting that simply adding bicycles cannot resolve unsatisfied demand.

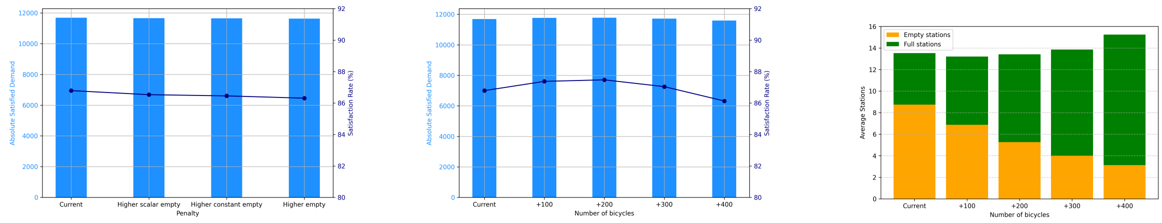

(a) Penalties **(b)** Satisfaction with extra bicycles **(c)** Empty and full stations with extra bicycles

Figure 9.7: Analysis of Empty Stations for Single-Type System

Figure 9.8 shows the analysis of empty stations for the mixed-fleet system. The additional bicycles considered here are electric, since earlier results showed that stations were short of electric bicycles more frequently. Electric bicycles also have higher usage rates. Increasing the scalar for nearly empty stations, raising the constant for empty stations, or doing both does not improve the satisfaction rate. In contrast, adding extra bicycles can increase satisfaction to over 98%, satisfying almost all demand. However, adding more than 200 bicycles does not lead to further improvement, and adding 300 or 400 bicycles actually reduces satisfaction because it increases the number of full stations. The optimal addition is 200 electric bicycles, which increases satisfaction above 98% and reduces the number of critical stations (empty + full) by more than half. This confirms that the availability of electric bicycles is a key bottleneck in mixed-fleet system of Valladolid.

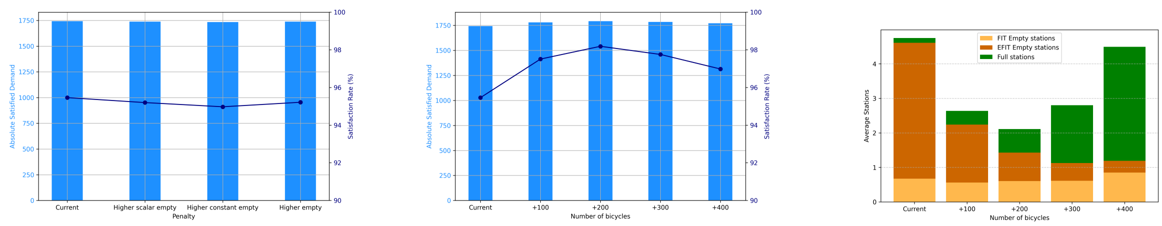

(a) Penalties **(b)** Satisfaction with extra bicycles **(c)** Empty and full stations with extra bicycles

Figure 9.8: Analysis of Empty Stations for Mixed-Fleet System

10 Conclusion

This thesis presented a dynamic repositioning strategy for continuously operating shared bicycle systems, applicable to single-type and mixed-fleet bicycle configurations. The primary objective was to improve user satisfaction while maintaining operational efficiency through flexible, time-dependent planning horizons. User satisfaction was quantified as a no-service penalty, reflecting (near) empty and full station states and deviations from optimal inventory levels. Operational efficiency was enforced through a maximum route duration constraint, which adapted to the time of day with shorter durations during peak hours, and longer durations during off-peak and nighttime periods.

For the single-type system in Zaragoza, the no-repositioning strategy had the lowest performance, meeting just over 80% of the demand. Static repositioning offered modest improvements, while the dynamic strategy substantially increased satisfied demand by approximately 400 additional trips daily. Even with perfect knowledge of future demand, more than 10% of the trips remained unserved and an average of 7.5 stations remained empty. For the mixed-fleet system in Valladolid, similar trends were observed. The static strategy added around 80 trips on top of the no-repositioning strategy. The dynamic strategy further improved performance, satisfying more than 95% of total demand. In line with the single-type system, demand shortages were concentrated at the trip origins, particularly for electric bicycles, reflecting the imbalance between the high demand for electric bicycles and limited availability. Further analysis revealed that adding 200 extra bicycles improves satisfaction in both systems, and in Valladolid it allows nearly all demand to be satisfied. Adjusting key parameters showed that the number of vehicles affects satisfaction levels, while vehicle size has little to no impact. Similarly, vehicle-aware repositioning has little effect on performance, though it can help reduce computation time. Lastly, both systems are capable of handling an increase in demand of 50% while experiencing only a decrease of 4% in satisfaction level.

In summary, this thesis demonstrates the successful development of a dynamic repositioning strategy with flexible, time-dependent planning horizon that substantially increases satisfied demand, and consequently, user satisfaction. This strategy makes shared bicycle systems more attractive as a sustainable alternative to the use of cars and public transport. The results show that the approach is effective even in systems with high usage, such as in Zaragoza, where four million trips are made annually, highlighting its potential to support widespread adoption and long-term impact.

11
Discussion

The solution strategy presented in this thesis is able to obtain desirable outcomes, but also have some limitations. To improve the suggested methods, these limitations can be further analyzed and the method can be extended. In this thesis, travel and stop times are treated as deterministic, whereas in practice they are uncertain. Incorporating stochasticity would make the model more robust to variability in these durations. In the simulations, users are assumed neither to switch to another bicycle type when their preferred one is unavailable nor to go to a different station to start their trip. Incorporating probabilities for these behaviors would make the simulation tool more realistic. Future research could focus on developing targeted strategies for specific bottlenecks in the system. When critical stations and time periods are identified, it would be possible to apply specific solutions, such as more frequent visits or adjusted prioritization rules, to enhance overall system performance. Another valuable extension would be to incorporate broken bicycles into the repositioning model. These bicycles take up docking space without contributing to availability, thereby reducing service levels. The efficient return of broken bicycles to the depot not only frees up space but also allows repairs, ultimately increasing the number of functional bicycles in the system (Wang & Szeto, 2018). In addition to charging stations, a bicycle sharing system may also contain solar-powered stations. These solar-powered stations present unique operational challenges. Solar-powered stations do not charge bicycles. Solar energy is used exclusively to provide availability data. As a result, bicycles with low battery levels parked at these stations cannot be used by riders and occupy dock space, similar to broken bicycles. Managing the relocation of these low-charge bicycles to charging stations adds another layer of complexity, requiring strategic decisions that impact both bicycle availability and dock space utilization.

References

Alvarez-Valdes, R., Belenguer, J. M., Benavent, E., Bermudez, J. D., Muñoz, F., Vercher, E., & Verdejo, F. (2016). Optimizing the level of service quality of a bike-sharing system. *Omega*, *62*, 163–175.

Angelelli, E., Chiari, M., Mor, A., & Speranza, M. G. (2022). A simulation framework for a station-based bike-sharing system. *Computers & Industrial Engineering*, *171*, 108489.

Bieliński, T., Kwapisz, A., & Ważna, A. (2021). Electric bike-sharing services mode substitution for driving, public transit, and cycling. *Transportation research part D: transport and environment*, *96*, 102883.

Brinkmann, J., Ulmer, M. W., & Mattfeld, D. C. (2015). Short-term strategies for stochastic inventory routing in bike sharing systems. *Transportation Research Procedia*, *10*, 364–373.

Brinkmann, J., Ulmer, M. W., & Mattfeld, D. C. (2019). Dynamic lookahead policies for stochastic-dynamic inventory routing in bike sharing systems. *Computers & Operations Research*, *106*, 260–279.

Caggiani, L., Camporeale, R., Ottomanelli, M., & Szeto, W. Y. (2018). A modeling framework for the dynamic management of free-floating bike-sharing systems. *Transportation Research Part C: Emerging Technologies*, *87*, 159–182.

Caggiani, L., & Ottomanelli, M. (2013). A dynamic simulation based model for optimal fleet repositioning in bike-sharing systems. *Procedia-Social and Behavioral Sciences*, *87*, 203–210.

Campbell, A. A., Cherry, C. R., Ryerson, M. S., & Yang, X. (2016). Factors influencing the choice of shared bicycles and shared electric bikes in beijing. *Transportation research part C: emerging technologies*, *67*, 399–414.

Chemla, D., Meunier, F., & Calvo, R. W. (2013). Bike sharing systems: Solving the static rebalancing problem. *Discrete Optimization*, *10*(2), 120–146.

DeMaio, P. (2009). Bike-sharing: History, impacts, models of provision, and future. *Journal of public transportation*, *12*(4), 41–56.

Galatoulas, N.-F., Genikomsakis, K. N., & Ioakimidis, C. S. (2020). Spatio-temporal trends of e-bike sharing system deployment: A review in europe, north america and asia. *Sustainability*, *12*(11), 4611.

Gast, N., Massonnet, G., Reijsbergen, D., & Tribastone, M. (2015). Probabilistic forecasts of bike-sharing systems for journey planning. *Proceedings of the 24th ACM international on conference on information and knowledge management*, 703–712.

Ghamami, M., & Shojaei, M. (2018). Introducing a design framework for a multi-modal public transportation system, focusing on mixed-fleet bike-sharing systems. *Transportation Research Record*, *2672*(36), 103–115.

Hernández-Pérez, H., & Salazar-González, J.-J. (2004). A branch-and-cut algorithm for a traveling salesman problem with pickup and delivery. *Discrete Applied Mathematics*, *145*(1), 126–139.

Jiménez-Meroño, E., & Soriguera, F. (2024). Optimization of bike-sharing repositioning operations: A reactive real-time approach. *EURO Journal on Transportation and Logistics*, *13*, 100138.

Kloimüllner, C., Papazek, P., Hu, B., & Raidl, G. R. (2014). Balancing bicycle sharing systems: An approach for the dynamic case. *European conference on evolutionary computation in combinatorial optimization*, 73–84.

Legros, B. (2019). Dynamic repositioning strategy in a bike-sharing system; how to prioritize and how to rebalance a bike station. *European Journal of Operational Research*, *272*(2), 740–753.

Li, Y., Szeto, W. Y., Long, J., & Shui, C. S. (2016). A multiple type bike repositioning problem. *Transportation Research Part B: Methodological*, *90*, 263–278.

Martins Silva, M. C., Aloise, D., & Jena, S. D. (2023). Towards effective rebalancing of bike-sharing systems with regular and electric bikes.

O'Brien, O. (2025). Bike sharing world map. Retrieved March 13, 2025, from https://bikesharingworldmap.com

O'Mahony, E., & Shmoys, D. (2015). Data analysis and optimization for (citi) bike sharing. *Proceedings of the AAAI conference on artificial intelligence*, *29*(1).

Pfrommer, J., Warrington, J., Schildbach, G., & Morari, M. (2014). Dynamic vehicle redistribution and online price incentives in shared mobility systems. *IEEE Transactions on Intelligent Transportation Systems*, *15*(4), 1567–1578.

Rainer-Harbach, M., Papazek, P., Hu, B., & Raidl, G. R. (2013). Balancing bicycle sharing systems: A variable neighborhood search approach. *Evolutionary Computation in Combinatorial Optimization: 13th European Conference, EvoCOP 2013, Vienna, Austria, April 3-5, 2013. Proceedings 13*, 121–132.

Raviv, T., & Kolka, O. (2013). Optimal inventory management of a bike-sharing station. *Iie Transactions*, *45*(10), 1077–1093.

Raviv, T., Tzur, M., & Forma, I. A. (2013). Static repositioning in a bike-sharing system: Models and solution approaches. *EURO Journal on Transportation and Logistics*, *2*(3), 187–229.

Regue, R., & Recker, W. (2014). Proactive vehicle routing with inferred demand to solve the bikesharing rebalancing problem. *Transportation Research Part E: Logistics and Transportation Review*, *72*, 192–209.

Salazar-González, J.-J., & Santos-Hernández, B. (2015). The split-demand one-commodity pickup-and-delivery travelling salesman problem. *Transportation Research Part B: Methodological*, *75*, 58–73.

Schuijbroek, J., Hampshire, R. C., & Van Hoeve, W.-J. (2017). Inventory rebalancing and vehicle routing in bike sharing systems. *European Journal of Operational Research*, *257*(3), 992–1004.

Shui, C. S., & Szeto, W. Y. (2018). Dynamic green bike repositioning problem–a hybrid rolling horizon artificial bee colony algorithm approach. *Transportation Research Part D: Transport and Environment*, *60*, 119–136.

Shui, C. S., & Szeto, W. Y. (2020). A review of bicycle-sharing service planning problems. *Transportation Research Part C: Emerging Technologies*, *117*, 102648.

Soriguera, F., & Jiménez-Meroño, E. (2020). A continuous approximation model for the optimal design of public bike-sharing systems. *Sustainable Cities and Society*, *52*, 101826.

Wang, Y.-J., Kuo, Y.-H., Huang, G. Q., Gu, W., & Hu, Y. (2022). Dynamic demand-driven bike station clustering. *Transportation Research Part E: Logistics and Transportation Review*, *160*, 102656.

Wang, Y., & Szeto, W. Y. (2018). Static green repositioning in bike sharing systems with broken bikes. *Transportation Research Part D: Transport and Environment*, *65*, 438–457.

Zhang, D., Yu, C., Desai, J., Lau, H., & Srivathsan, S. (2017). A time-space network flow approach to dynamic repositioning in bicycle sharing systems. *Transportation research part B: methodological*, *103*, 188–207.

Zhu, S. (2021). Optimal fleet deployment strategy: Model the effect of shared e-bikes on bike-sharing system. *Journal of Advanced Transportation*, *2021*(1), 6678637.

A

Scientific paper

Dynamic Shared Bicycle Repositioning with Flexible Time Horizon in a Mixed Fleet System

Tom Bruinink[a,1,*]

[a]*TU Delft, Stevinweg 1, 2628 CN, Deft, The Netherlands*

Abstract

Shared bicycle systems offer a sustainable alternative to car and public transport use in urban areas, but maintaining a balance between bicycle availability and demand remains a challenge. This thesis presents a time-adaptive dynamic repositioning strategy that jointly optimizes station selection, bicycle quantities, and routing for continuously operating systems with a mixed bicycle fleet. User satisfaction is captured through a no-service penalty, while operational efficiency is ensured by enforcing time-dependent maximum route durations, shorter during peak hours and longer during off-peak and at night. Applied to a real-world case study with a mixed bicycle fleet in Valladolid, the strategy significantly increases satisfied demand and user satisfaction. These improvements enhance the potential of the system to promote sustainable urban mobility on a broader scale.

Keywords: bike-sharing, dynamic repositioning, station-based, mixed-fleet, MILP, simulation

1. Introduction

Shared bicycle systems have become a key element of sustainable urban transportation by reducing greenhouse gas emissions, travel time, and traffic congestion, while improving connectivity to other modes of transit by addressing the first-mile/last-mile challenge (DeMaio, 2009). Over the last two decades, the number of shared bicycle systems has expanded rapidly, with more than 2000 systems and nearly 10 million bicycles operating worldwide (O'Brien, 2025). However, this rapid growth has introduced new operational challenges. With an increasing number of stations and users, maintaining a balance between supply and demand has become more complex. The imbalance in bicycle availability between different areas throughout the day reduces system efficiency and user satisfaction, ultimately limiting system impact and growth potential. To address this, operators deploy repositioning vehicles to redistribute bicycles. At the same time, the adoption of mixed bicycle fleets that combine electric and mechanical bicycles, further enhances the attractiveness of the system by improving accessibility, allowing longer trips and reducing physical effort for users (Ghamami and Shojaei, 2018). However, this diversification also introduces additional complexity to the repositioning problem, as a surplus of one bicycle type cannot necessarily compensate for a shortage of another, which requires more sophisticated decision-making to balance both types across stations.

The most common strategy to address the bicycle repositioning problem is a static approach, where bicycles are redistributed during off-peak hours, mainly at night (Alvarez-Valdes et al., 2016; Chemla et al., 2013). Using this approach, repositioning routes are planned to prepare the system for the expected demand the following day. However, this method struggles to respond to real-time system dynamics and heavily relies on demand predictions, which are often uncertain. Furthermore, preventing imbalances during the day would require an excessively large fleet of bicycles, making this approach inefficient for dynamic environments.

An alternative strategy is the dynamic approach, in which repositioning occurs during operation, based on the real-time system state and short-term demand forecasts. Repositioning vehicles are deployed on short routes that serve only the most critical stations. This method offers several advantages: It can respond to fluctuations in demand, depends less on demand forecasts, and has a low computational cost, since typically only one vehicle is optimized at a time, specifically the one that has just completed its previous route.

In this paper, we propose an integrated dynamic repositioning model that simultaneously optimizes station selection, bicycle movements, and routing in real time. The model incorporates a flexible planning horizon, which is shorter during peak hours and longer during off-peak periods, to balance service quality and operational efficiency. To ensure scalability and coordination in multi-vehicle operations, the model uses vehicle-aware repositioning, taking into account the expected end location of all vehicle to guide future task assignments. The inclusion of a mixed bicycle fleet further increases the complexity of the bicycle repositioning problem, as repositioning decisions must consider both bicycle types to maintain a balanced availability which enhances service quality for diverse user needs.

The contributions of this paper can be summarized as follows:

- Operational efficiency and user satisfaction measures are defined for continuously operating shared bicycle systems.

*Corresponding author

- A time-adaptive dynamic bicycle repositioning problem is proposed that integrates station selection, bicycle quantity decisions, and efficient routing planning for continuously operating shared bicycle systems with mixed bicycle fleet.
- Valuable insights are provided on the growing importance of bicycle repositioning, based on realistic case study results.

This paper is structured as follows. Section 2 reviews the existing literature, while Section 3 introduces the problem. The model is presented in Section 4 and Section 5 explains the simulation method. Section 6 discusses the case study and Section 7 presents the results. Finally, Section 8 and Section 9 summarize the main conclusions and describe some future research directions.

2. Literature review

We first introduce the dynamic approach and then discuss the different approaches to improve it. An overview of the main characteristics of the literature reviewed in this section is provided in Table 1. Only papers explicitly addressing the bicycle repositioning problem are included; studies without routing or relocation considerations are excluded due to their fundamentally different scope and modeling assumptions.

2.1. Dynamic approach

The goal of the dynamic approach to bicycle repositioning is to maintain balanced bicycle availability across stations in real time, ensuring high user satisfaction by minimizing empty and full stations during operations. The dynamic approach leverages real-time station data and short-term demand predictions to make adaptive, responsive decisions throughout the day. A key feature of this method is that new repositioning routes are generated only after the previous route is completed. This sequential decision-making process significantly reduces computational complexity, as only one route is constructed at a time, rather than solving an entire network of repositioning actions simultaneously (Jiménez-Meroño and Soriguera, 2024). The primary objective of dynamic repositioning is to minimize unsatisfied demand by proactively adjusting station inventories before they become empty or full. Although factors such as minimizing travel time or aligning station inventories with target fill levels are still considered, they are generally secondary to service quality (Jiménez-Meroño and Soriguera, 2024). Studies have shown that even minimal repositioning efforts can lead to substantial improvements in service levels (Legros, 2019). Interestingly, the optimal inventory level for each station often varies throughout the day based on demand fluctuations, but tends to be independent of vehicle capacity (Legros, 2019). To support these decisions, various prioritization policies have been developed to select which stations to visit. These include targeting stations with high current or future unmet demand, large deviations from desired inventory levels, high movement frequency, or proximity to other stations (Legros, 2019). In contrast, stations with consistently low demand are deprioritized to avoid inefficient repositioning operations (Jiménez-Meroño and Soriguera, 2024).

Various approaches have been developed to solve the bicycle repositioning problem, differing in how they determine which stations to serve, the number of bicycles to relocate, and the routes to follow. One approach to address the bicycle repositioning problem is the action-based approach, which simplifies decision-making by focusing on the next station to visit and the number of bicycles to pick up or deliver. By focusing on individual actions rather than full routes, this method greatly reduces computational complexity. It enables fast responsiveness to real-time station conditions and short-term demand forecasts, ensuring that the most critical stations are addressed quickly (Jiménez-Meroño and Soriguera, 2024; Legros, 2019). For example, Jiménez-Meroño and Soriguera (2024) and Legros (2019) model the problem as a pairwise task assignment, where a repositioning vehicle is continuously assigned to a single action. Likewise, Brinkmann et al. (2015) introduce a framework that distinguishes between short-term and long-term strategies. Their short-term approach prioritizes nearby stations with immediate violation risks, while the long-term approach targets the most imbalanced stations overall. However, a common limitation of these action-based strategies is that decisions are made sequentially, considering only one step at a time. This can lead to inefficiencies as opportunities to combine actions or optimize across multiple stations are missed. For example, visiting two nearby unbalanced stations together or allowing simultaneous pickup and delivery actions might produce better results than handling each action in isolation (Legros, 2019).

Another approach is the two-step strategy, which breaks the repositioning problem into sequential decision phases. Pfrommer et al. (2014) first identify promising repositioning routes and then determine the optimal number of bicycles to move along each route, ultimately selecting the route with the highest utility. In contrast, Regue and Recker (2014) first optimize the station inventory levels and then construct a route based on the optimal repositioning quantities. To reduce computational complexity, they only consider nearby stations with significant inventory imbalances. Although two-step methods can improve computational efficiency and impose structure, they are limited by the strong influence of first-stage decisions. Fixing part of the solution early may reduce flexibility in later stages, potentially leading to suboptimal outcomes.

The integrated approach to dynamic bicycle repositioning combines the decisions of determining which stations to visit, the optimal number of bicycles to relocate, and the most efficient route (Kloimüllner et al., 2014). By addressing all of these subproblems together in a single model, the integrated approach enables vehicles to receive optimized repositioning routes, ultimately improving both operational efficiency and user satisfaction.

An important dimension in dynamic bicycle repositioning is the time horizon over which decisions are planned and updated. Two common formulations are the rolling-horizon and the time–space network approaches, which can be applied to different repositioning strategies regardless of whether they are integrated, two-step, or action-based. The rolling-horizon method divides the operational period into a sequence of fixed-duration intervals, each treated as a static subproblem (Shui and

Table 1: Classification of literature on dynamic bicycle repositioning approaches

Reference	Strategy	Optimization method	Time horizon	Demand predictions	Fleet type	Coordination	Main objective
Caggiani and Ottomanelli (2013)	Integrated	Simulation-based	Rolling horizon	Short-term	Single	No coordination	Hybrid
Kloimüllner et al. (2014)	Integrated	Heuristic	Rolling horizon	Short- & Long-term	Single	No coordination	User satisfaction
Pfrommer et al. (2014)	Two-step	Heuristic	Rolling horizon	Short-term	Single	Vehicle-aware	User satisfaction
Regue and Recker (2014)	Two-step	Exact	Rolling horizon	Short-term	Single	No coordination	Hybrid
O'Mahony and Shmoys (2015)	Integrated	Exact/Heuristic	Time-adaptive	Short- & Long-term	Single	Clustering	User satisfaction
Brinkmann et al. (2015)	Action-based	Simulation-based	Rolling horizon	-	Single	No coordination	User satisfaction
Zhang et al. (2017)	Two-step	Exact/Heuristic	Time-space	Short-term	Single	No coordination	Hybrid
Shui and Szeto (2018)	Two-step	Heuristic	Rolling horizon	Short-term	Single	No coordination	Hybrid
Caggiani et al. (2018)	Integrated	Simulation-based	Rolling horizon	Short-term	Single	Clustering	User satisfaction
Brinkmann et al. (2019)	Action-based	Simulation-based	Time-adaptive	Short- & Long-term	Single	No coordination	User satisfaction
Legros (2019)	Action-based	Exact	Rolling horizon	Short-term	Single	Clustering	User satisfaction
Jiménez-Meroño and Soriguera (2024)	Action-based	Heuristic	Rolling horizon	Short-term	Single	Vehicle-aware	User satisfaction
This study	Integrated	Exact	Time-adaptive	Short-term	Mixed	Vehicle-aware	User satisfaction

Szeto, 2018). During each stage, real-time station states and short-term demand forecasts are assumed to be known, effectively making the problem static for that interval. After executing the repositioning routes, the system is updated with new inventory levels and demand estimates before the next stage begins. This method balances adaptability with simplified optimization at each step. A common implementation of this method is to divide the day into five-minute intervals (Caggiani and Ottomanelli, 2013; Caggiani et al., 2018; Shui and Szeto, 2020). This short interval lengths provides sufficient opportunities to react to fluctuations in demand while keeping the computational effort manageable. The time–space network approach instead models operations on a continuous-time basis, providing a more detailed representation of system dynamics. This method, as explored by Zhang et al. (2017), facilitates making real-time decisions by capturing the evolving nature of both the repositioning process and station demand. A key advantage of this approach is its ability to take into account the actual station status at the time the repositioning vehicle arrives, resulting in more accurate repositioning actions. However, flexibility in modeling comes with a significant increase in computational complexity. To manage this, Zhang et al. (2017) proposed a decomposition into a two-stage optimization model, reformulating the original nonlinear problem into a mixed-integer program. Although this makes the problem more tractable, such a staged approach inherently leads to suboptimal solutions, as decisions in the first stage do not fully anticipate outcomes of the second. As a result, this approach has seen limited application in the literature and remains impractical for large-scale, real-world systems.

2.2. Model extensions

To align with the developments in bicycle sharing systems, multiple model extensions will be considered. These extensions are independent of the repositioning strategy.

Most existing dynamic repositioning models assume homogeneous fleets, which limits their applicability to real-world systems that often offer a variety of bicycle types, such as electric and mechanical bicycles (Shui and Szeto, 2020). The model proposed by Li et al. (2016) introduces this heterogeneity considering bicycles with varying characteristics. Zhu (2021) studies how the integration of electric bicycles into a mechanical fleet impacts system performance. This introduces new operational challenges: A surplus of one bicycle type cannot necessarily compensate for a shortage of another, and different bicycle types may occupy varying amounts of space in both repositioning vehicles and docking stations. To address this, Martins Silva et al. (2023) determine target inventory levels for each bicycle type and examine how the willingness of users to substitute their preferred bicycle type influences system performance. These complexities highlight the need for more refined loading and unloading strategies, as well as more advanced decision-making processes to ensure that the correct type of bicycle is delivered to meet the specific user demand.

Another extension is the time-adaptive approach, considered as a variant of the rolling-horizon method, which accounts for the heterogeneous demand patterns throughout the day, recognizing that the operator's priorities typically align with these varying demand levels. O'Mahony and Shmoys (2015) propose using a clustering model during the day and a routing model overnight, acknowledging that different periods require tailored strategies. Brinkmann et al. (2019) propose a policy that simulates future demand over a predefined horizon, with decision-making horizons that vary throughout the day to capture heterogeneous demand patterns. These time-dependent look-ahead horizons are autonomously set using the value function approximation, enabling the system to anticipate peak-hour commuter demand without extending simulations unnecessarily. From an operational perspective, their model assumes that repositioning vehicles operate continuously throughout the day and that costs related to routing and drivers are already incurred. As a result, transportation costs are excluded from the optimization, allowing the model to focus primarily on maintaining high service levels.

A final extension aimed at reducing computational complexity and improving vehicle coordination is vehicle-aware repositioning, which incorporates the expected end locations of vehicles into the decision-making process. Traditionally, many studies address complexity by clustering stations into different areas, each assigned to a repositioning vehicle. For example, Chemla et al. (2013) and Legros (2019) divide the city into exclusive zones to avoid coordination issues and allow independent vehicle routing. Similarly, O'Mahony and Shmoys (2015) propose clustering stations during peak hours to manage high demand and system fluctuations. Clustering helps to reduce the problem size and avoids long truck routes. Although clustering offers clear computational advantages, it can reduce flexibility in dynamic environments. To address this, a more adaptive strategy, vehicle-aware repositioning, directly incorporates the locations of other vehicles into the repositioning model. For example, Pfrommer et al. (2014) and Jiménez-Meroño and Soriguera (2024) explicitly consider the expected end locations of all vehicles in the system when assigning a new task. This look-ahead approach retains some of the efficiency benefits of clustering while improving flexibility, as decisions are based on both current and anticipated vehicle positions.

3. Problem definition

To evaluate the trade-off between operational efficiency and user satisfaction, we consider a daily operation of a shared bicycle system. The system must function across different time periods, overnight, peak, and off-peak hours, each characterized by different demand patterns. Throughout the day, multiple repositioning vehicles continuously relocate bicycles to maintain a balanced inventory between stations, based on the current number of bicycles at each station and the expected demand in the upcoming period. Each driver completes one task at a time, usually involving multiple actions, such as picking up or delivering bicycles. After completion of a task, the driver requests a new one. Therefore, for each vehicle, a route must be created by assigning the next task to be performed, without knowing which tasks will follow. We adopt a rolling-horizon approach in which the time horizon is discretized into minutes. The planning horizon for each task is set equal to the maximum route duration applicable to that time period, with routes allowed to depart and return at any minute. We begin by presenting the dynamic framework, which outlines the steps involved in creating the next route, starting with the driver's request for the next task and ending with the execution of that task. Thereafter we explain briefly the real-time station status, the demand predictions, and the no-service penalties.

3.1. Dynamic framework

The steps of creating the next route are presented in the dynamic framework (Figure 1). The process begins in the upper-left corner, where a driver requests a new task after completing the previous one. This request enters the platform, which analyzes the real-time status of all stations defined as the number of available bicycles and free docking spots. Additionally, demand predictions are generated for the upcoming period. The system status and demand predictions are used as inputs to the no-service level function. This value is computed for each station using the function $f(S, D)$, where S represents the current inventory level and D denotes the predicted demand. Together, these values represent the expected inventory at each station if no-repositioning actions are performed. The resulting no-service penalties are then used as inputs for the Dynamic Bicycle Repositioning Problem with Flexible Time Horizon (DBRPFTH). This model determines which stations to visit, how many bicycles to relocate, and the most efficient route to follow. A detailed description of the model is provided in Section 4. Once the next route is determined, the task is assigned to the driver for execution. Upon completion, the driver submits a new task request, and the cycle repeats.

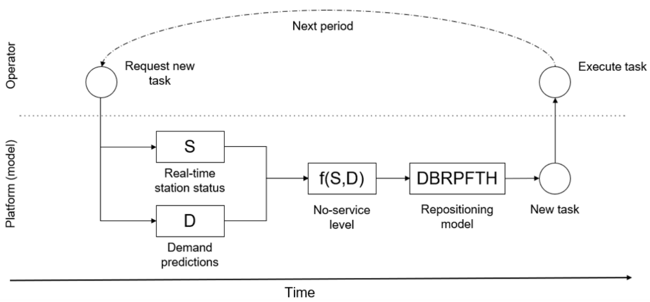

Figure 1: Schematic overview of the dynamic framework

3.2. Real-time station status

The real-time station status describes the current state at each station. It contains the number of bicycles available and the number of free docking spots. Data will be collected from open-access data sources in the standard General Bikeshare Feed Specification (GBFS) format, which contains real-time information on shared bicycle systems. It is updated every 30 seconds.

3.3. Demand predictions

Demand predictions estimate fluctuations in bicycle levels at each station for the next time interval. To ensure flexibility, the prediction horizon is adapted to the planning horizon of each repositioning route. During peak hours, shorter forecasts of 30 minutes are preferred due to rapid changes in station inventory levels. In contrast, longer prediction windows of 40 to 45 minutes are more suitable for overnight and off-peak periods when fluctuations are less frequent. In addition, the predictions are aligned dynamically with route schedules. For example, if a route ends at 10:40 with a half-hour forecast horizon, predictions are generated for 10:40–11:10. This is achieved by partitioning the predictions into five-minute intervals, allowing flexible adaptation to different start times and planning horizons.

The prediction model relies solely on unconstrained historical demand data, which means that changes in bicycle inventory are unaffected by station capacity limits. If a station has no bicycles and is still empty in the next period, the actual demand might be nonzero, but any demand is unobservable. Similarly, when a station reaches full capacity, the number of users who wanted to return a bicycle but were unable to do so remains unknown. To ensure accuracy, we only use historical data where the station inventory remained within capacity limits during the next interval.

The demand prediction model consists of several steps. The first step involves clustering past dates based on their demand patterns, without considering calendar characteristics. The number of trips for all stations will be aggregated to a total demand for a day. Note that if the demand pattern for several stations differs significantly from the usual pattern, the day will be in a different cluster. Importantly, rainy days are excluded from clustering because the rain negatively affects demand, potentially leading to misclassifications.

In the second step, a machine learning (ML) classification model predicts the cluster to which a future date belongs based on its calendar characteristics, such as the day of the week, month, holidays, and special events.

In the third step, an ML regression model is used to forecast demand for the next time interval based on multiple input features. These include the assigned cluster for the given date, lagged demand values from previous time periods, demand fluctuations observed on similar past dates within the same cluster, calendar variables, and the most recent weather forecasts, including rainfall, temperature, wind speed, and humidity.

To improve reliability, two quantile regression models are trained to estimate the lower and upper demand bounds. These quantile models are trained with loss functions that emphasize penalties on either negative demand (more bicycles taken than returned) or positive demand (more bicycles returned than taken), which also an estimation of the 20th and 80th percentiles, respectively. To improve robustness against demand uncertainty, the system adopts a worst-case scenario approach based on the estimated demand bounds. When the lower and upper quantiles indicate that two additional bicycles might be picked up or dropped off and the station is currently empty, the scenario with more pickups is considered the worst case. In contrast, if the station is full, the scenario with more deliveries is considered worse. This conservative approach helps mitigate the risk of demand fluctuations, ensuring a more reliable and stable system.

3.4. No-service penalties

User satisfaction will be quantified using no-service penalties. They are determined using real-time station status and demand predictions, as these two factors allow us to determine the expected inventory level for each station in the upcoming period, assuming no repositioning actions are taken. The ideal scenario for each station is a balance between the number of bicycles and free docks, ensuring that the station is neither empty nor full and is capable of accommodating future demand. Since future demand is uncertain, stations that are near-empty or near-full are considered suboptimal. A no-service penalty will be imposed if the station occupancy becomes nearly empty or nearly full.

The problem definition outlined in this section forms the basis for the mathematical formulation presented in the next section.

4. Model formulation

In this section, we introduce the Dynamic Bicycle Repositioning Problem with Flexible Time Horizon (DBRPFTH), which focuses on determining the stations to visit, the optimal number of bicycles to relocate, and creating an efficient route with a flexible time horizon for a mixed bicycle fleet. We formulate the DBRPFTH as a standard Mixed Integer Linear Program (MILP) to optimize user satisfaction while ensuring operational efficiency. This model is based on the Static Bicycle Repositioning Problem (SBRP) introduced by Raviv et al. (2013). We begin by introducing the measures for operational efficiency and user satisfaction before delving deeper into the model.

4.1. Measures

Operational efficiency is important in evaluating repositioning strategies with the primary goal of creating efficient routes that respond effectively to fluctuations in station inventory levels. In our context, where the number of drivers is predetermined and vehicles operate continuously, traditional cost-minimization objectives, such as reducing driver wages and vehicle operations, are less relevant. Since drivers request new routes immediately after completing the previous one, minimizing routing costs does not significantly impact efficiency. Instead, we focus on creating efficient routes to maximize user satisfaction.

We introduce a maximum route duration T, which determines how long a vehicle can operate before requesting a new route. The route duration is influenced by the travel time between stations, which varies throughout the day, and the stop time, which includes both a fixed time for the pickup or delivery of bicycles and an additional component based on the number of bicycles being moved. To account for variations in travel time and system conditions throughout the day, we divide the day into three distinct time intervals: peak hours, off-peak hours, and overnight. Each time period falls within one of these intervals, allowing parameters to be defined at the interval level rather than for each individual time period. To further adapt to these variations, we use different duration matrices for travel times in each interval, ensuring that travel times accurately reflect the conditions during repositioning routes.

In order to address variations in demand and station inventory fluctuations throughout the day, we introduce the time-dependent maximum route duration, T_t, where t represents the specific time period. During peak hours, when demand is high and station inventory fluctuates significantly, it becomes crucial to respond quickly to these fluctuations by planning shorter routes. To achieve this, we reduce the maximum route duration T_t during peak hours, ensuring that repositioning routes can adapt faster to changes in inventory levels. On the other hand, during nighttime and off-peak hours, when demand is lower and station inventory levels are more stable, the model shifts its focus towards optimizing operational efficiency. In these periods, longer routes can be planned because the station inventory fluctuates less, allowing more actions to be efficiently included in a single route. To accommodate this, the maximum route duration T_t is extended, which allows the system to optimize for fewer, but longer, repositioning tasks.

User satisfaction is assessed by no-service penalties, as described in Subsection 3.4. A penalty is imposed if the expected station occupancy, determined by the current inventory level and the demand forecast for the upcoming period, falls below a lower threshold (nearly empty) or exceeds an upper threshold (nearly full). These penalties are calculated based on the squared deviation from the thresholds, which places greater emphasis on larger deviations. This reflects the fact that extreme imbalances can result in unmet demand. Note that we have a mixed-fleet. An excess of one bicycle type may not effectively resolve the shortages of another (Shui and Szeto, 2020). Moreover, the impact of repositioning a single bicycle can go beyond the benefit of a single user, as it may enable a chain of successful rentals and returns (Chemla et al., 2013). In addition, a fixed penalty is applied if a station becomes completely empty or full. Together, these mechanisms discourage extreme inventory levels and promote more balanced stations, ultimately aiming to improve service levels and reduce the risk of long-term negative consequences.

To determine the optimal inventory level for each station in the upcoming periods, demand forecasts are incorporated for both the upcoming period and the period after. The demand in the upcoming period reflects the demand during which repositioning takes place, while the forecast for the two periods ahead represents the expected demand after repositioning has occurred. For example, if a station is expected to become (nearly) full in the coming period, we know that a pickup is needed, but it remains unclear whether this should involve removing 25%, 50%, or 75% of the bicycles. The forecast for the period after repositioning helps refine this decision: If demand is expected to be positive (more returns than withdrawals), it is preferable to leave fewer bicycles at the station; if demand is expected to be negative (more withdrawals than returns), having more bicycles available is advantageous. This component is integrated into a quadratic penalty function, where the penalty is lowest at the target inventory level that best aligns with anticipated demand. A relatively small weight is assigned to this term to ensure that inventory levels outside the acceptable range are still penalized more heavily than those within it. This structure prioritizes meeting immediate demand predictions while using future forecasts to fine-tune inventory within the preferred bounds.

4.2. Notation

The dynamic repositioning problem is described by the following set and parameters:

N	Set of nodes, including the depot, indexed by $i = 0, \ldots, \lvert N \rvert$.
i_o	Origin node, where the repositioning operation starts.
s_i^{0m}	Number of mechanical bicycles at node i before the repositioning operation starts.
s_i^{0e}	Number of electric bicycles at node i before the repositioning operation starts.
d_i^{1m}	Predicted demand of mechanical bicycles at node i for the next period.
d_i^{1e}	Predicted demand of electric bicycles at node i for the next period.
d_i^{2m}	Predicted demand of mechanical bicycles at node i for two periods ahead.
d_i^{2e}	Predicted demand of electric bicycles at node i for two periods ahead.
c_i	Number of docks installed at node i, referred to as the node's capacity.
c^{max}	Maximum number of docks.
l_m	Initial load of mechanical bicycles in the vehicle.
l_e	Initial load of electric bicycles in the vehicle.
k	Capacity (number of bicycles) of the vehicle.
$f_i(s_i^m, s_i^e)$	Penalty function reflecting user satisfaction for node i.
r_{ij}	Traveling time from node i to node j.
T_t	Repositioning time, i.e., maximum duration of a repositioning route in period t.
V	Constant time required for visiting a node.
P	Time required to pick up a bicycle from a node.
D	Time required to deliver a bicycle to a node.
M	Upper bound on number of arcs, default is $\lvert N \rvert$

Note that to reduce computational complexity only imbalanced stations are considered (Regue and Recker, 2014). The depot is assumed to have no demand and capacity, because the vehicle cannot drop-off or pick up bicycles at the depot. Since the driver requests a new route after completing the previous one, the vehicle does not necessarily start its route from the depot. Instead, the starting location is referred to as the origin node, denoted by i_o. Station capacity is defined as the total number of usable bicycles and docks, excluding any that are broken or unavailable. Additionally, we assume that all docks are available for both types of bicycles, that each bicycle type occupies the same amount of space in the repositioning vehicle, and that the time required to pick up or deliver a mechanical or electric bicycle is identical.

The following decision variables will be used:

x_{ij}	Binary variable which equals one if the vehicle travels directly from node i to node j, and zero otherwise.
y_{ij}^m	Number of mechanical bicycles carried on the vehicle when it travels directly from node i to node j. If the vehicle does not travel directly from i to j, y_{ij}^m equals zero.
y_{ij}^e	Number of electric bicycles carried on the vehicle when it travels directly from node i to node j. If the vehicle does not travel directly from i to j, y_{ij}^e equals zero.
y_i^{Pm}	Number of mechanical bicycles picked up at node i.
y_i^{Pe}	Number of electric bicycles picked up at node i.
y_i^{Dm}	Number of mechanical bicycles delivered at node i.
y_i^{De}	Number of electric bicycles delivered at node i.
z_i^{Pm}	Binary variable which equals one if mechanical bicycles are picked up from node i, and zero otherwise.
z_i^{Pe}	Binary variable which equals one if electric bicycles are picked up from node i, and zero otherwise.
z_i^{Dm}	Binary variable which equals one if mechanical bicycles are delivered to node i, and zero otherwise.
z_i^{De}	Binary variable which equals one if electric bicycles are delivered to node i, and zero otherwise.
z_i	Binary variable which equals one if action is performed at node i, and zero otherwise.
q_i	Auxiliary variable used for sub-tour elimination constraints.
s_i^m	Expected inventory level of mechanical bicycles at node i at the end of the repositioning operation.
s_i^e	Expected inventory level of electric bicycles at node i at the end of the repositioning operation.

4.3. Objective function

Our objective is to maximize user satisfaction by minimizing no-service penalties, using the satisfaction measures defined in Subsection 4.1. This approach is supported by Caggiani and Ottomanelli (2013), who model utility as a plateau function, where utility increases from an empty station to a sufficient lower bound, remains high within an optimal range, and decreases as the station nears full capacity. However, in contrast to their approach, we aim to determine the optimal inventory level within the acceptable range (the plateau) based on the demand predictions for the period after repositioning, which corresponds to two periods ahead.

We introduce the penalty function to evaluate the expected inventory levels. This function combines a convex component and a piecewise-linear component to reflect the penalties related to imbalanced stations, based on the research of Raviv et al. (2013). To incorporate this function into an optimization model, we reformulate both components in linear terms using auxiliary variables and constraints. This transformation enables the complete model to be expressed as a Mixed-Integer Linear Program (MILP), following the approach of Raviv et al. (2013). It is important to note that the values of $f_i(s_i^m, s_i^e)$ can be calculated from the input variables.

The function contains the following parameters:

L_m	Percentage of station capacity considered as almost empty of mechanical bicycles
L_e	Percentage of station capacity considered as almost empty of electric bicycles
H	Percentage of station capacity considered as almost full
p^L	Penalty (almost) empty station
p^H	Penalty (almost) full station
p^{cL}	Constant penalty empty station
p^{cH}	Constant penalty full station
β	Weight/scaling factor for optimal inventory level

The penalty function for station $i \in N$ given station inventory level s_i^m, s_i^e is as follows:

$$\begin{aligned} f_i(s_i^m, s_i^e) = {} & p^L(\max(c_i \cdot L_m - s_i^m, 0))^2 + p^L(\max(c_i \cdot L_e - s_i^e, 0))^2 \\ & + p^H(\max(s_i^m + s_i^e - c_i \cdot H, 0))^2 \\ & + \beta \cdot \Big[(s_i^m - (\frac{L_m + 0.5 \cdot H}{2} \cdot c_i - d_i^{2m}))^2 \\ & + (s_i^e - (\frac{L_e + 0.5 \cdot H}{2} \cdot c_i - d_i^{2e}))^2\Big] \\ & + p^{cL} 1_{s_i^m \le 0} + p^{cL} 1_{s_i^e \le 0} + p^{cH} 1_{s_i^m + s_i^e \ge c_i} \quad \forall i \in N \end{aligned} \tag{1}$$

The penalty function consists of eight components, which can be divided into two types: convex components and piecewise linear components. The first five components form the convex part of the function, while the last three components introduce piecewise linear penalties. Consequently, the penalty function $f_i(s_i^m, s_i^e)$ can be separated into a convex penalty function $f_i^C(s_i^m, s_i^e)$ and a piecewise linear penalty function $f_i^L(s_i^m, s_i^e)$.

$$\begin{aligned} f_i^C(s_i^m, s_i^e) = {} & p^L(\max(c_i \cdot L_m - s_i^m, 0))^2 + p^L(\max(c_i \cdot L_e - s_i^e, 0))^2 \\ & + p^H(\max(s_i^m + s_i^e - c_i \cdot H, 0))^2 \\ & + \beta \cdot \Big[(s_i^m - (\frac{L_m + 0.5 \cdot H}{2} \cdot c_i - d_i^{2m}))^2 \\ & + (s_i^e - (\frac{L_e + 0.5 \cdot H}{2} \cdot c_i - d_i^{2e}))^2\Big] \quad \forall i \in N \\ f_i^L(s_i^m, s_i^e) = {} & p^{cL} 1_{s_i^m \le 0} + p^{cL} 1_{s_i^e \le 0} + p^{cH} 1_{s_i^m, s_i^e \ge c_i} \quad \forall i \in N \end{aligned}$$

The first and second components of the convex penalty function address situations where a station's inventory level s_i^m or s_i^e drops below a low threshold, indicating that the station is nearly empty of mechanical or electric bicycles. This penalty is squared and weighted by the scalar p^L, increasing the penalty as the inventory level is further from the threshold. The third component focuses on stations that are (nearly) full, applying a similar squared penalty with a scalar p^H as the total station inventory level $s_i^m + s_i^e$ approaches the station's full capacity.
The fourth and fifth components are always active, but are specifically used in determining the optimal inventory level within the acceptable range. It distinguishes between inventory levels that are still balanced, between the lower and upper bounds, by using the demand forecast for two periods ahead, which is the demand in the period after repositioning. If the post-repositioning demand forecast is zero, the optimal inventory level for each bicycle type lies in the middle of its lower bound and half of the overall upper bound. Combined, the optimal total inventory level of both types lies in the middle of the sum of their lower bounds and the overall upper bound. If the forecast is positive (indicating more returns than withdrawals), it is better to leave fewer bicycles at the station to prevent overflow. As a result the penalty function then favors lower inventory levels. In contrast, if the forecast is negative (more withdrawals than returns), the system benefits from higher bicycle availability, and the penalty is minimized at higher inventory levels. This preference is modeled using a quadratic penalty function, where the minimum penalty corresponds to the inventory level that best matches the expected post-repositioning demand. A small scalar weight β is applied to this term to ensure that inventory levels outside the boundaries are always worse than inventory levels within the boundaries. This ensures that the model prioritizes meeting immediate demand (during repositioning). Future forecasts are used to fine-tune inventory levels within the bounds. This results in more balanced and forward-looking repositioning decisions.
Piecewise linear components introduce additional penalties when a station becomes either completely empty or completely full. These components are captured by an indicator function that equals 1 if s_i^m or s_i^e is non-positive (station is empty) or greater than or equal to the capacity of the station c_i (station is full). A negative inventory level is possible when the station expects a large demand deficit, meaning that significantly more bicycles are taken than returned. Similarly, an inventory level exceeding the station's capacity can occur when the station expects a large demand surplus, meaning that significantly more bicycles are returned than taken.
Together, the above mentioned eight components ensure that the repositioning strategy not only addresses immediate demand and balances stations within acceptable inventory levels but also accounts for future demand fluctuations, minimizes the likelihood of stations becoming entirely empty or full, and prioritizes critical stations that directly impact service levels. This multi-component penalty structure enables the model to make forward-looking, balanced repositioning decisions that improve both operational efficiency and user satisfaction.

To transform the convex penalty function f_i^C into linear terms, we must first determine the domain of the expected inventory levels s_i^m and s_i^e. This level depends on the initial inventories before repositioning $s_i^{0_m}$ and $s_i^{0_e}$, the predicted demands $d_i^{1_m}$ and $d_i^{1_e}$, and the repositioning actions $y_i^{P_m}$, $y_i^{P_e}$ (pickups), $y_i^{D_m}$ and $y_i^{D_e}$ (drop-offs). The initial inventories $s_i^{0_m}$ and $s_i^{0_e}$ range from 0 to station capacity c_i, while $d_i^{1_m}$ and $d_i^{1_e}$ can take any integer value. Repositioning actions $y_i^{P_m}$, $y_i^{P_e}$, $y_i^{D_m}$ and $y_i^{D_e}$ are limited by the number of bicycles at the station or the free docks available before repositioning. If $d_i^{1_m}$ and $d_i^{1_e}$ are zero, the inventories s_i^m and s_i^e will remain within the range $[0, c_i]$. However, if $d_i^{1_m}$ or $d_i^{1_e}$ is negative, which means that more bicycles are taken than returned, s_i^m or s_i^e can drop below zero, extending the lower bound of the domain to $d_i^{1_m}$ or $d_i^{1_e}$. Conversely, if $d_i^{1_m}$ or $d_i^{1_e}$ is positive, which means that more bicycles are returned than taken, s_i^m or s_i^e can exceed the station's capacity, extending the upper bound to $d_i^{1_m} + c_i$ or $d_i^{1_e} + c_i$. Thus, the overall domains of s_i^m and s_i^e are given by:

$$s_i^m \in [min(0, d_i^{1m}), max(c_i, c_i + d_i^{1m})]$$
$$s_i^e \in [min(0, d_i^{1e}), max(c_i, c_i + d_i^{1e})]$$

The domains for the first two and last two components of the convex penalty function f_i^C are s_i^m and s_i^e respectively. For the third component the domains are combined, since the upper bound depends on the total inventory level. To replace the convex penalty functions f_i^C with a linear term and linear constraints, we introduce the following sets and equations:

$$U_1 = [min(0, d_i^{1m}), max(c_i - 1, c_i - 1 + d_i^{1m})]$$
$$U_2 = [min(0, d_i^{1e}), max(c_i - 1, c_i - 1 + d_i^{1e}]$$
$$U_3 = [min(0, d_i^{1m} + d_i^{1e}), max(c_i - 1, c_i - 1 + d_i^{1m} + d_i^{1e})]$$
$$b_{iu} \equiv f_i^C(u+1) - f_i^C(u) \qquad \forall i \in N, u \in U_1, U_2, U_3$$
$$a_{iu} \equiv f_i^C(u) - b_{iu} \cdot u \qquad \forall i \in N, u \in U_1, U_2, U_3$$

The sets U_1, U_2 and U_3 define the index ranges over which s_i^m, s_i^e and $s_i^m + s_i^e$ are defined, respectively. The parameter b_{iu} captures the marginal penalty associated with adding the $(u+1)^{th}$ bicycle to the station i. Together, a_{iu} and b_{iu} represent the intercept and slope, respectively, of the linear function that approximates the convex penalty function f_i^C at the u^{th} level. These linear approximations are included as constraints in the MILP to ensure that the correct value of the convex part is considered in the objective function, depending on the expected inventory levels s_i^m and s_i^e. The resulting value of the convex penalty function for each station is captured by the variables $g_i^{L_m}$ (lower bound mechanical), $g_i^{L_e}$ (lower bound electric), g_i^H (upper bound), $g_i^{IB_m}$ (optimal level mechanical) and $g_i^{IB_e}$ (optimal level electric), which are included in the objective function.

To linearize the piecewise penalty function f_i^L, we replace the indicator functions in $f^L(\cdot)$ with binary decision variables. Specifically, we define binary variables w_i^m, w_i^e and o_i, which indicate whether station i is empty or full, respectively. That is, $w_i^m = 1$ if the inventory level $s_i^m \leq 0$, $w_i^e = 1$ if the inventory level $s_i^e \leq 0$ and $o_i = 1$ if the inventory level $s_i^m + s_i^e \geq c_i$. To enforce this behavior, we introduce three linear constraints using a sufficiently large constant C.

4.4. Mathematical model

The mathematical model for the dynamic bicycle repositioning problem is as follows:

$$\min \quad \sum_{i \in N} g_i^{Lm} + g_i^{Le} + g_i^H + g_i^{IBm} + g_i^{IBe} + p^{cL} \cdot (w_i^m + w_i^e) + p^{cH} \cdot o_i \tag{2}$$

$$\text{s.t.} \quad g_i^{Lm} \geq a_{iu_1}^{Lm} + b_{iu_1}^{Lm} \cdot s_i^m \qquad \forall i \in N, u_1 \in U_1 \tag{3}$$
$$g_i^{Le} \geq a_{iu_2}^{Le} + b_{iu_2}^{Le} \cdot s_i^e \qquad \forall i \in N, u_2 \in U_2 \tag{4}$$
$$g_i^H \geq a_{iu_3}^H + b_{iu_3}^H \cdot (s_i^m + s_i^e) \qquad \forall i \in N, u_3 \in U_3 \tag{5}$$
$$g_i^{IBm} \geq a_{iu_1}^{IBm} + b_{iu_1}^{IBm} \cdot (s_i^m) \qquad \forall i \in N, u_1 \in U_1 \tag{6}$$
$$g_i^{IBe} \geq a_{iu_2}^{IBe} + b_{iu_2}^{IBe} \cdot (s_i^e) \qquad \forall i \in N, u_2 \in U_2 \tag{7}$$
$$s_i^m > -C \cdot w_i^m \qquad \forall i \in N \tag{8}$$
$$s_i^e > -C \cdot w_i^e \qquad \forall i \in N \tag{9}$$
$$s_i^m + s_i^e < c_i + C \cdot o_i \qquad \forall i \in N \tag{10}$$
$$s_i^m = s_i^{0m} + d_i^{1m} - (y_i^{Pm} - y_i^{Dm}) \qquad \forall i \in N \tag{11}$$
$$s_i^e = s_i^{0e} + d_i^{1e} - (y_i^{Pe} - y_i^{De}) \qquad \forall i \in N \tag{12}$$
$$y_i^{Pm} - y_i^{Dm} = \sum_{j \in N, j \neq i} y_{ij}^m - \sum_{j \in N, j \neq i} y_{ji}^m \qquad \forall i \in N \setminus i_o \tag{13}$$
$$y_i^{Pe} - y_i^{De} = \sum_{j \in N, j \neq i} y_{ij}^e - \sum_{j \in N, j \neq i} y_{ji}^e \qquad \forall i \in N \setminus i_o \tag{14}$$
$$l_m + y_{i_o}^{Pm} - y_{i_o}^{Dm} = \sum_{j \in N, j \neq i_o} y_{i_o j}^m \tag{15}$$
$$l_e + y_{i_o}^{Pe} - y_{i_o}^{De} = \sum_{j \in N, j \neq i_o} y_{i_o j}^e \tag{16}$$
$$y_{ij}^m + y_{ij}^e \leq k \cdot x_{ij} \qquad \forall i, j \in N, i \neq j \tag{17}$$
$$\sum_{j \in N, j \neq i} x_{ij} = \sum_{j \in N, j \neq i} x_{ji} \qquad \forall i \in N \tag{18}$$
$$y_i^{Pm} \leq s_i^{0m} \qquad \forall i \in N \tag{19}$$
$$y_i^{Pe} \leq s_i^{0e} \qquad \forall i \in N \tag{20}$$
$$y_i^{Dm} + y_i^{De} \leq c_i - (s_i^{0m} + s_i^{0e}) \qquad \forall i \in N \tag{21}$$
$$\sum_{i \in N} \Big(\sum_{j \in N} r_{ij} \cdot x_{ij} + P \cdot (y_i^{Pm} + y_i^{Pe}) + D \cdot (y_i^{Dm} + y_i^{De}) + V \cdot z_i \Big) \leq T_t \tag{22}$$
$$q_j \geq q_i + 1 - M(1 - x_{ij}) \qquad \forall i \in N, j \in N \setminus i_o, i \neq j \tag{23}$$
$$\frac{y_i^{Pm}}{c^{max}} \leq z_i^{Pm} \leq y_i^{Pm} \qquad \forall i \in N \tag{24}$$
$$\frac{y_i^{Pe}}{c^{max}} \leq z_i^{Pe} \leq y_i^{Pe} \qquad \forall i \in N \tag{25}$$
$$\frac{y_i^{Dm}}{c^{max}} \leq z_i^{Dm} \leq y_i^{Dm} \qquad \forall i \in N \tag{26}$$
$$\frac{y_i^{De}}{c^{max}} \leq z_i^{De} \leq y_i^{De} \qquad \forall i \in N \tag{27}$$
$$z_i^{Pm} + z_i^{Dm} <= 1 \qquad \forall i \in N \tag{28}$$
$$z_i^{Pe} + z_i^{De} <= 1 \qquad \forall i \in N \tag{29}$$
$$\frac{z_i^{Pm} + z_i^{Pe} + z_i^{Dm} + z_i^{De}}{4} \leq z_i \leq z_i^{Pm} + z_i^{Pe} + z_i^{Dm} + z_i^{De} \qquad \forall i \in N \tag{30}$$
$$x_{ij} \in \mathbb{B} \qquad \forall i, j \in N \tag{31}$$
$$z_i^{Pm}, z_i^{Pe}, z_i^{Dm}, z_i^{De}, z_i \in \mathbb{B} \qquad \forall i \in N \tag{32}$$
$$y_i^{Pm}, y_i^{Pe}, y_i^{Dm}, y_i^{De} \in \mathbb{Z}_{\geq 0} \qquad \forall i \in N \tag{33}$$
$$y_{ij}^m, y_{ij}^e \geq 0 \qquad \forall i, j \in N, i \neq j \tag{34}$$
$$q_i \geq 0 \qquad \forall i \in N \tag{35}$$

The objective function (2) minimizes the no-service penalties, as described by (1). Constraints (3)-(7) determine the convex no-service penalties based on the linearization intercept and slope, while constraints (8)-(10) set the piecewise-linear no-service penalties. The inventory balance is enforced by constraints (11) and (12), which set the expected inventory levels after repositioning. Furthermore, the bicycle flow is conserved through constraints (13) and (14), while constraints (15) and (16) ensure that the flow leaving the origin node equals the initial vehicle load plus any actions taken at the origin. The constraints (17) limit the number of bicycles carried by a vehicle to its capacity and ensure that no bicycles are transported along an arc if the vehicle does not use that arc.

Vehicle flow is preserved through constraints (18), while the number of bicycles picked up or delivered to a station is restricted by constraints (19)-(21), ensuring that the pickups do not exceed the available inventory and that deliveries do not exceed the remaining docking capacity. Constraints (19) and (20) inherently enforce non-negativity of inventory variables, while constraints (21) ensure that the inventory at each station and the depot remains within capacity limits. As a result, explicit constraints for inventory non-negativity and capacity limits are not required.

The maximum route duration is constrained by (22). Note that the return to the origin node is set to zero to enforce a round-trip structure in the model, while the actual route will not return. The sub-tour elimination is enforced through constraints (23), which prevent small loops in the route. Constraints (24)-(27) ensure that an action is only performed if bicycles are picked up or delivered. Furthermore, constraints (28) and (29) ensure that a vehicle may only pick up or deliver a given bicycle type at a station, but not both for the same type. However, it may pick up one bicycle type and deliver another at the same station. Furthermore, constraints (30) ensure that stations are only visited when an action is performed.

Binary and integer constraints on decision variables are imposed by (31)-(33), while non-negativity conditions are enforced by constraints (34) and (35). Non-negativity constraints on s_i^m and s_i^e cannot be enforced, as the expected inventory level can be negative after repositioning when the demand deficit ($d_i^{1_m}$ and $d_i^{1_e}$) is highly negative. Similarly, s_i may exceed the station capacity c_i when demand is exceptionally high. In such cases, it is assumed that unmet demand is lost, but it is still taken into account in the objective function to quantify its impact. Finally, the integrality of y_{ij}^m, y_{ij}^e, s_i^m and s_i^e is implicitly ensured by the integer nature of $y_i^{P_m}$, $y_i^{P_e}$, $y_i^{D_m}$ $y_i^{D_e}$, $s_i^{0_m}$ and $s_i^{0_e}$.

Similar to Raviv et al. (2013) we introduce valid inequalities to reduce computation time.

$$\sum_{j \in N} x_{i_o j} \geq 1 \tag{36}$$

$$\sum_{j \in N, j \neq i} x_{ij} \leq 1 \qquad \forall i \in N \tag{37}$$

$$y_i^{Pm} \leq min(s_i^{0m}, k) \sum_{j \in N} x_{ij} \qquad \forall i \in N \tag{38}$$

$$y_i^{Pe} \leq min(s_i^{0e}, k) \sum_{j \in N} x_{ij} \qquad \forall i \in N \tag{39}$$

$$y_i^{Dm} + y_i^{De} \leq min(c_i - (s_i^{0m} + s_i^{0e}), k) \sum_{j \in N} x_{ij} \qquad \forall i \in N \tag{40}$$

$$y_i^{Pm} + y_i^{Pe} + y_i^{Dm} + y_i^{De} \geq \sum_{j \in N} x_{ij} \qquad \forall i \in N \setminus \{i_o, 0\} \tag{41}$$

Constraint (36) ensures that the vehicle departs from the origin at least once. In the LP relaxation, this departure count can take fractional values. By enforcing the departure of a complete vehicle, the vehicle flow-conservation constraints (18) guarantee that this condition holds across all visited nodes. Constraints (37) ensure that each station is visited at most once. Additionally constraints (38)-(40) impose a stricter form of constraints (19)-(21), limiting the number of bicycles picked up or delivered based on vehicle capacity and applying only when the vehicle visits the station. Note that in constraints (38) and (39), x_{ij} is 1 as well if bicycles of the other type are picked up. Finally, constraints (41) ensure that vehicles visit a station only if bicycles are picked up or delivered.

4.5. Vehicle-aware repositioning

Vehicle-aware repositioning is used to improve operational efficiency and reduce computational complexity while preserving vehicle flexibility without fixed spatial boundaries. This strategy incorporates the expected end locations of all repositioning vehicles directly into the decision-making process. This look-ahead approach enables the model to assign tasks in a way that anticipates future opportunities (Jiménez-Meroño and Soriguera, 2024).

In our approach, when a driver requests a new task, we explicitly consider the ongoing tasks of currently busy vehicles. If a busy vehicle is expected to finish its current task in a location close to a particular demand, that location is not assigned to the requesting driver. We limit potential task assignments by excluding a fixed number of locations near the expected end locations of busy vehicles. This vehicle-aware coordination improves fleet efficiency and scalability by flexibly limiting certain locations, without setting fixed area boundaries. Importantly, no such restrictions are applied to locations near the current position of the vehicle requesting the task.

Figure 2 shows how vehicle-aware repositioning affects the set of stations considered for assignment in Valladolid. In this example, four repositioning vehicles operate simultaneously. The green vehicle marks the position of the driver currently requesting a task. The three red vehicles represent the expected end locations of currently busy vehicles. Stations shown as green dots are the closest locations to the requesting vehicle, while red dots indicate stations located near the expected end points of busy vehicles. These stations are temporarily restricted for the requesting vehicle. Blue dots represent all other stations that remain available. The visual representation shows how certain areas of the network are effectively "reserved" for other vehicles that are better positioned to serve them in the near future.

In the scenario with ten restricted locations, some overlap occurs between green and red dots, meaning certain stations are both close to the requesting vehicle and within a restricted area. Proximity is calculated using travel time rather than straight-line (Euclidean) distance, so stations that appear geographically close on the map may still be farther away in practice. The model also excludes stations already assigned to other vehicles, stations previously visited by the requesting vehicle, and stations that are currently balanced. Consequently, the number of restricted locations directly influences the set of stations considered in the optimization.

5. Simulation

To evaluate the performance of bicycle repositioning strategies, we will perform a series of simulation experiments comparing four operational approaches: (1) no repositioning, (2) a static repositioning strategy, (3) the proposed state-of-the-art dynamic repositioning strategy developed in this study and (4) a

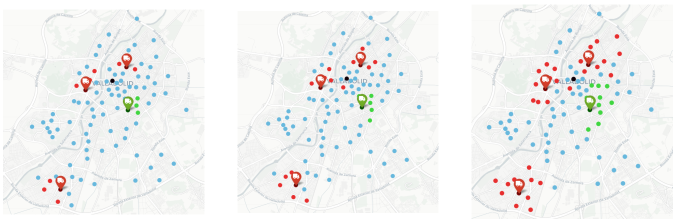

(a) 3 restricted locations (b) 5 restricted locations (c) 10 restricted locations

Figure 2: Vehicle-aware repositioning

perfect information strategy. In the static repositioning strategy, bicycles are redistributed only during the night shift, when demand is relatively low. The purpose of this strategy is to prepare the system for the following day. This can be seen as a simplified variant of the proposed dynamic repositioning approach, restricted to nighttime operations with a larger number of repositioning vehicles. As a result, the expected demand during the day is not explicitly considered, only the demand that occurs during the repositioning period itself is taken into account.

The perfect information approach serves as a benchmark by assuming that the dynamic repositioning strategy is executed with full knowledge of future demand. Instead of relying on forecasts, the actual demand for the upcoming planning horizon is provided to the strategy. This allows us to determine the maximum achievable satisfied demand under the proposed dynamic approach if the forecasts were perfectly accurate. Any remaining unsatisfied demand can then be attributed to resource limitations, such as the number or capacity of repositioning vehicles, available bicycles, or dock space, rather than to the repositioning strategy itself. This highlights the performance limits of the dynamic strategy and clarifies the extent to which lost demand is caused by factors other than forecast accuracy.

5.1. Demand

To replicate real-world system dynamics, user demand is simulated on the basis of historical station status data. Origin–destination (OD) trips are generated by Poisson sampling, which models discrete trip events and ensures system stability.

The expected demand for each OD pair in a given hour is calculated as the average demand from historical data. Similar to demand predictions, periods with constrained data must be excluded from historical data. Such constrained periods occur when stations are empty or full, making the actual demand unobservable during these times. A departure station is only considered constrained when it is empty (since a full station does not restrict departures), while an arrival station is constrained only when it is full. The demand for these constrained periods is determined based on the level of available unconstrained data. If more than half of the hour consists of unconstrained periods, the constrained demand is estimated proportionally, assuming that demand would have been linear relative to the observed unconstrained periods. However, if more than half of the hour is constrained, the remaining unconstrained data is considered unreliable. In such cases, the demand is replaced by the average hourly demand for that specific OD pair.

Since many OD pair combinations are rare events, either because stops are located very close to each other or too far apart, there are many OD pairs for which no demand was observed in the historical data. As a result, the OD matrix of the expected demand contains many zeros. This would imply that the Poisson-distributed demand for these OD pairs is always zero, preventing the simulation of any trips along these pairs. However, such rare trips do occasionally occur and therefore a small probability should be assigned. To address this issue, Laplace smoothing is applied. This technique assigns a small positive value to all OD pairs to prevent zero probabilities. To ensure that the total expected demand remains unchanged, the probabilities are subsequently rescaled. Specifically, for each OD pair (i, j), the smoothed probability is calculated as:

$$P_{ij}^{smoothed} = \frac{\lambda_{ij} + \gamma}{\lambda_{total} + \gamma V}$$

where:

- λ_{ij} is the original expected demand for OD pair (i, j)
- λ_{total} is the total expected demand over all OD pairs
- V is the total number of OD pairs
- γ is the smoothing parameter

The smoothed expected demand for OD pair (i, j) is then scaled back to ensure that the total demand remains unchanged:

$$\hat{\lambda}_{ij} = \lambda_{total} \cdot P_{ij}^{smoothed} = \lambda_{total} \cdot \frac{\lambda_{ij} + \gamma}{\lambda_{total} + \gamma V}$$

This approach guarantees that the sum of all smoothed demands remains equal to the original total demand, while allowing all OD pairs to remain possible in simulations.

5.2. Experiment

The simulation incorporates three repositioning shifts per day, each lasting eight hours: a night shift (22:00–06:00), a morning shift (06:00–14:00), and an afternoon shift (14:00–22:00). To ensure smoother operations, the start and end times of some vehicles within a shift may be slightly adjusted. For example, the start of a shift for part of the fleet may be delayed by an hour to avoid all vehicles being at the depot simultaneously. Each vehicle starts its shift by leaving the depot and returns there at the end of the shift. During operations, drivers request a new task after completing their previous one. The next task can include stations that are unbalanced, not visited by other vehicles and not visited in the previous route of the vehicle. The repositioning algorithm then determines the next assignment. If no feasible route is found at that moment, the driver will wait for five minutes before requesting a new task. The system is updated every minute, taking into account both user trips and repositioning movements. When a bicycle trip and a repositioning action occur at the same station within the same minute, the repositioning trip is prioritized. If no bicycles of the requested type are available at the departure station, the user cannot start

the trip, and the request is recorded as unsatisfied demand. An important assumption is that users do not switch bicycle types if their preferred type is unavailable. This implies that a shortage of one bicycle type cannot be compensated by a surplus of another. If no available docks are found at the intended destination station, the user parks the bicycle at the nearest station with available docks. However, this is still counted as unsatisfied demand at the intended destination station. Note that when the simulation day begins, there may already be bicycles moving, trips that started before the simulation's start time but arrive after it has begun. To account for these trips, historical trip data is used to represent bicycles that were actually moving at the time the system status is retrieved. These are not simulated movements but real trips that started before the simulation period and are scheduled to arrive after its start, ensuring the initial system status reflects the true state of bicycles on the network. When a vehicle arrives at a station to pick up or deliver bicycles, but the required number of bicycles (for pickup) or free docks (for delivery) is not available, the vehicle adjusts by picking up or delivering as many as possible. This limitation may affect the feasibility of subsequent planned actions, as the number of bicycles on the vehicle may differ from what was initially intended. If the remaining bicycle stock on the vehicle does not allow the original plan to continue, the planned quantities for the following stations are adjusted accordingly or the remaining actions are canceled.
The standard analysis will focus on a typical weekday under normal conditions. In addition, the simulation will investigate the effects of several factors. Each simulation scenario will be assessed using the validation metrics introduced in Section 7. This comprehensive approach will provide insights in the performance of different repositioning strategies under various operating conditions.

6. Case Study

We evaluate the performance of our models by means of simulation experiments using data from a station-based bicycle-sharing system in Valladolid. Data are collected from open-access data sources with standard General Bikeshare Feed Specification (GBFS) formatted data, which contain real-time information for shared bicycle operations. The predictions and simulated demand used in our models are generated as described in Subsection 3.3 and Section 5.
The shared bicycles system in Valladolid serves around 1,800 daily users and operates a mixed fleet of 420 mechanical and 415 electric bicycles distributed over 100 stations. The system offers a total of 1846 docks, resulting in a docks-to-bicycles ratio of 2.2, which is favorable for efficient operation (Soriguera and Jiménez-Meroño, 2020). The daily trips-per-bicycle ratio is relatively low at 2.2, indicating low usage compared to systems with higher demand. The average number of docks per station is 18.5, providing sufficient capacity to accommodate returning bicycles and reduce the risk of full stations.
Table 2 presents the parameter settings. The system operates with four vehicles during the day and one at night. In the static approach, four vehicles are used during the night to prepare the system for the upcoming day. Vehicle capacities range from 11 to 14 bicycles, with no restrictions on locations visited. Maximum route durations depend on the time period: 45 minutes at night, 40 minutes off-peak, and 30 minutes during peak hours. The first task of each shift has a maximum route duration of 45 minutes to allow the vehicle to depart from the depot. If the remaining shift time is less than the maximum route duration plus the time needed to return to the depot, the vehicle must return to the depot to ensure it finishes its shift on time. Each stop includes a fixed handling time plus an additional time per bicycle moved. Penalty values are consistent across scenarios, with lower inventory bounds set per vehicle type and the upper bound applied to total station occupancy. The scalar p^L for (near) empty stations is higher than p^H for (near) full stations, as empty stations are considered more problematic. Denied departures often mean users leave the system, while in the case of full stations, users continue their trip by diverging to nearby stations. The smoothing parameter for demand estimation ensures that the expected demand ratio between OD pairs with and without observed trips is at least 50.

Table 2: Parameter settings

Category	Parameter	Value
Number of vehicles	Number of vehicles (night)	1
	Number of vehicles (morning)	4
	Number of vehicles (afternoon)	4
Vehicle restrictions	Vehicle capacity	11-14
	Restricted locations	0
Duration (min)	Maximum duration (night)	45
	Maximum duration (off-peak)	40
	Maximum duration (peak)	30
Stop-time (min)	Constant stop-time	5
	Stop-time per bicycle	0.5
Penalties	β	0.001
	Lower bound	0.15
	Upper bound	0.7
	pL	1
	p^H	0.5
	p^{cL}	20
	p^{cH}	20
Simulation	γ	0.0002

7. Results

The results are obtained by applying the simulation methods as explained in Section 5 to the model explained in Section 4 to the case study introduced in Section 6. All optimization problems are solved using Gurobi 12.0.1, with an average computation time of approximately 5 seconds per task. For each scenario, 10 simulation runs are performed to obtain stable results, and the results are presented as the average values along with their standard deviations.
The main KPI to evaluate the performance of bicycle repositioning strategies is satisfied demand, which reflects the number of users able to find an available bicycle at their origin and a free dock at their destination. This measure directly represents user satisfaction, the main objective in shared bicycle operations. In contrast, unsatisfied demand at origins (due to empty stations) and destinations (due to full stations) highlights potential imbalances. Unlike satisfied demand, the no-service penalty is not suitable for comparing repositioning strategies, as it is zero

Table 3: Average results

Type	No repositioning	Static	Dynamic	Perfect information
Demand satisfaction				
Satisfied demand	1542.6	1622.8	1743.5	1782.7
Unsatisfied demand origin	250.2	178	79.6	42.7
Unsatisfied demand destination	36.2	27.5	3.3	0.6
Satisfaction rate (%)	84.3	88.8	95.5	97.6
Station status				
Empty stations mechanical bicycles	5.6	2.6	0.7	0.3
Empty stations electric bicycles	13.1	9	3.9	2.4
Full stations	1.8	0.7	0.1	0
Number of mechanical bicycles at stations	415.1	412.5	411.6	410.8
Number of electric bicycles at stations	400.2	396.8	391.1	391.5
Repositioning				
Number of tasks	0	44	115.1	116
Number of stations visited	0	136.6	296.4	299.8
Number of bicycles repositioned	0	180.9	451.4	474.3
Route duration (s)	-	2570.2	2197.2	2180

for the no-repositioning strategy and generally lower for the static strategy due to night-time operations. Additional validation metrics include the number of empty and full stations, the number of bicycles at stations, the total number of repositioning tasks, stations visited and bicycles moved, and route durations.

7.1. General results

The average results for mixed-fleet system in Valladolid are presented in Table 3. As expected, the scenario without repositioning results in the lowest level of satisfied demand, with the system serving around 1540 trips. The static repositioning strategy already improves this by approximately 80 additional trips, while the dynamic strategy achieves a substantial further increase, satisfying 120 more users and covering over 95% of total demand. With perfect information about future demand, the system could only serve about 40 additional trips (2%) compared to the regular dynamic strategy, suggesting that the dynamic approach already comes close to the maximum achievable performance.

The unsatisfied demand at the trip origins is notably higher than at the destinations for all strategies. This indicates that the primary reason for unsatisfied demand is the lack of available bicycles rather than the lack of dock space. Notably, for the dynamic strategy and perfect information scenarios the unsatisfied demand at the destinations is reduced to almost zero. The distinction between mechanical and electric bicycles further reveals that stations experience electric empty states more frequently than mechanical empty states, highlighting the greater challenge of maintaining the availability of electric bicycles. With the dynamic approach, an average of 3.9 stations have no electric bicycles available compared to only 0.7 stations lacking mechanical bicycles. This disparity aligns with the considerably higher demand for electric bicycles, which is more than 4 times higher than the demand for mechanical bicycles. The number of mechanical and electric bicycles at stations is similar across the different strategies. The difference between strategies is slightly larger for electric bicycles (around 10) than for mechanical bicycles (around 5). This is consistent with the higher usage of electric bicycles, which probably also increases the need for their repositioning.

With regard to repositioning movements, the static strategy leads to more visited stations per task compared to the dynamic strategy, probably because the allowed route duration is longer during the night period. The number of bicycles moved per station is relatively low for all strategies, indicating many small-scale repositioning moves. The total number of tasks, stations visited, and bicycles moved in the regular dynamic approach and the dynamic approach with perfect information is very similar, implying that further improvements in demand prediction could increase satisfied demand without additional operational effort. Finally, the average route duration for the static strategy is considerably higher than for the dynamic and perfect information strategies, as expected, since it takes place only during the night. The dynamic and perfect information strategies yield nearly identical route durations, indicating that route length is not a differentiating factor between these approaches.

To analyze the repositioning strategies in more detail, we first look at the demand pattern and the number of bicycles at stations throughout the day (Figure 3). These patterns are very similar across all strategies, with only subtle differences. Therefore we present them only for the dynamic approach. Demand stays relatively high until midnight, after which it drops sharply to nearly zero. Around 6:00 AM, demand begins to rise again, reaching a morning peak between 7:00 and 9:00 AM, then declining significantly until about 1:00 PM. The highest demand peak occurs between 2:00 and 3:00 PM, after which demand remains elevated until midnight. The results show that almost all unsatisfied demand is avoided, and repositioning helps stabilize the system following peak periods. Notably, unsatisfied demand at destination stations is barely visible. Furthermore, we observe that the confidence interval, represented by the shaded area around the mean, is narrow, indicating a high consistency and low variability in the results.

Figure 4 shows how station status changes over the day for

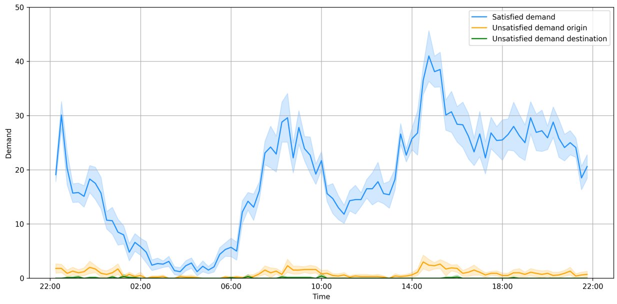

(a) Demand over the Day

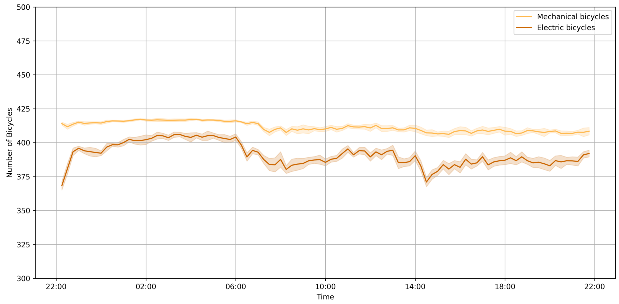

(b) Number of Bicycles at Stations over the Day

Figure 3: Analysis over the Day

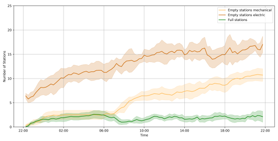

(a) No repositioning

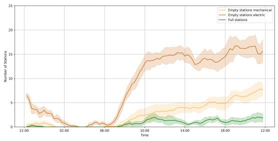

(b) Static repositioning

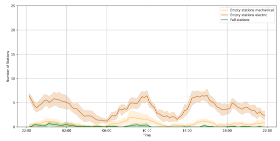

(c) Dynamic repositioning

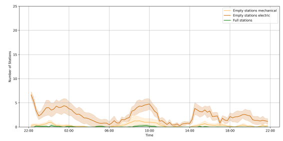

(d) Perfect information

Figure 4: Station Status over the Day

each repositioning strategy. There is a significant difference between the no-repositioning and static strategies on one side, and the dynamic approach (with perfect information) on the other. In the no-repositioning and static cases, where no repositioning is carried out during the day, the number of empty stations steadily rises, with particularly many stations lacking electric bicycles. Full stations are also present throughout the day. In contrast, the dynamic strategies keep the number of empty stations low during the day, with further reductions after peak periods and overnight. Full stations are almost negligible under these approaches. Although perfect information allows for slightly better handling after peak hours, its advantage over regular dynamic repositioning is relatively small.

7.2. Sensitivity analysis of system parameters

To gain a deeper understanding of how various model parameters affect system performance, an additional sensitivity analysis was performed on the dynamic repositioning strategy. First, we find that the initial system status has minimal impact on the satisfied demand, confirming the robustness of the results to variations in starting conditions. Figure 5 presents the results of the sensitivity analysis. The light blue bars represent the absolute number of satisfied demand, while the dark blue line indicates the satisfaction rate. It is worth noting that the absolute satisfied demand may increase even when the satisfaction rate decreases, and vice versa. Increasing the number of vehicles improves the satisfaction rate from 94% to 96%, with the largest gain observed when increasing from fewer vehicles to the current number. Varying vehicle size has a small effect on performance; even when decreasing vehicle capacity to just 9 bicycles, the satisfaction rate remains considerably high. The small differences observed are mainly due to randomness of demand. Incorporating demand predictions improves model performance, increasing the satisfaction rate by almost 1% and bringing it within 2% of the perfect information benchmark. Vehicle-aware repositioning lowers satisfaction slightly but reduces the computation time from 5 to 3 seconds. Its lower performance is likely due to the relatively small number of unbalanced stations, meaning that restricting the set of candidate stations greatly limits the solution space. Weekend and holidays share similar, significantly lower demand patterns. However, the satisfaction rate drops on weekends, probably due to the reduced repositioning fleet of only two vehicles during the day. Finally, when the system demand increases, the number of trips increases substantially, while the satisfaction rate decreases only moderately (by 1%–4%), remaining above 90%. This indicates that the system can handle significant demand growth with minimal loss in service quality.

7.3. Analysis of empty stations

In the results, we observed significantly more unsatisfied demand at the origin of a trip, more empty stations than full ones, and that even with perfect information it was difficult to avoid empty stations during the day. To investigate the root of this issue, we consider two options: increasing the penalties for empty stations to prioritize them, or adding bicycles to test whether a shortage of bicycles is the underlying cause. Figure 6 shows the analysis of empty stations. The additional bicycles considered here are electric, since earlier results showed that stations were short of electric bicycles more frequently. Electric bicycles also have higher usage rates. Increasing the scalar for nearly empty stations, raising the constant for empty stations, or doing both does not improve the satisfaction rate. In contrast, adding extra bicycles can increase satisfaction to over 98%, satisfying almost all demand. However, adding more than 200 bicycles does not lead to further improvement, and adding 300 or 400 bicycles actually reduces satisfaction because it increases the number of full stations. The optimal addition is 200 electric bicycles, which increases satisfaction above 98% and reduces the number of critical stations (empty + full) by more than half. This confirms that the availability of electric bicycles is a key bottleneck in mixed-fleet system of Valladolid.

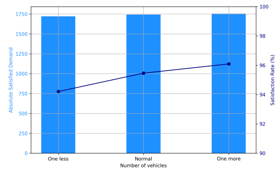

(a) Number of vehicles

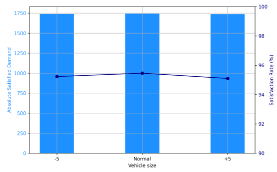

(b) Vehicle size

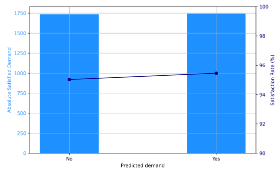

(c) Predicted demand

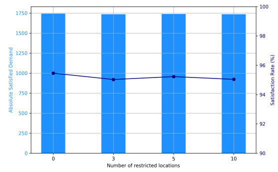

(d) Restricted vehicle locations

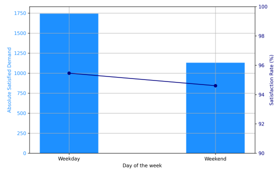

(e) Day of the week

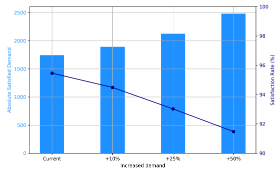

(f) Increased demand

Figure 5: Sensitivity Analysis of System Parameters

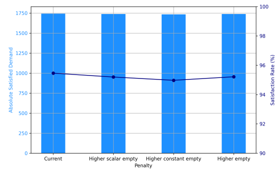

(a) Penalties

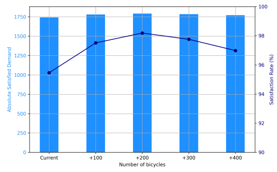

(b) Satisfaction with extra bicycles

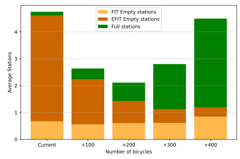

(c) Empty and full stations with extra bicycles

Figure 6: Analysis of Empty Stations for Mixed-Fleet System

8. Conclusion

This paper presented the successful implementation of a dynamic repositioning strategy for continuously operating shared bicycle systems with a mixed bicycle fleet. The primary objective was to improve user satisfaction while maintaining operational efficiency by flexible time-dependent planning horizons. The no-repositioning strategy had the lowest performance, whereas the static strategy offered modest improvements. The dynamic strategy further improved performance, satisfying more than 95% of total demand. Demand shortages were concentrated at the origin of a trip, particularly for electric bicycles, reflecting the imbalance between the high demand of electric bicycles and their limited availability. Further analysis revealed that adding 200 extra electric bicycles satisfies almost all demand. Adjusting key parameters showed that the number of vehicles affects satisfaction levels, while vehicle size has little to no impact. Similarly, vehicle-aware repositioning has little effect on performance, though it can help reduce computation time. Lastly, the system is capable of handling an increase in demand of 50% while the satisfaction level remains above 90%. In summary, this thesis demonstrates the successful development of a dynamic repositioning strategy with flexible, time-dependent planning horizon that substantially increases satisfied demand, and consequently, user satisfaction. The strategy makes shared bicycle systems more attractive as a sustainable alternative to the use of cars and public transport.

9. Further research directions

Future research could focus on developing targeted strategies for specific bottlenecks in the system. When these critical stations and time periods are identified, it would allow specific solutions, such as more frequent visits or adjusted prioritization rules, to enhance overall system performance. Another valuable extension would be to incorporate broken bicycles into the repositioning model. These bicycles take up docking space without contributing to availability, thereby reducing service levels. Efficiently returning broken bicycles to the depot not only frees up space but also allows repairs, ultimately increasing the number of functional bicycles in the system (Wang and Szeto, 2018). Additionally, solar-powered stations introduce unique operational challenges, because they do not charge bicycles. These stations operate on solar energy, which limits them to providing availability data only. As a result, bicycles with low battery levels parked at these stations cannot be used by riders and occupy dock space, similar to broken bicycles. Managing the relocation of these low-charge bicycles to charging stations adds another layer of complexity, requiring strategic decisions that impact both bicycle availability and dock space utilization.

References

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

Brinkmann, J., Ulmer, M.W., Mattfeld, D.C., 2015. Short-term strategies for stochastic inventory routing in bike sharing systems. Transportation Research Procedia 10, 364–373.

Brinkmann, J., Ulmer, M.W., Mattfeld, D.C., 2019. Dynamic lookahead policies for stochastic-dynamic inventory routing in bike sharing systems. Computers & Operations Research 106, 260–279.

Caggiani, L., Camporeale, R., Ottomanelli, M., Szeto, W.Y., 2018. A modeling framework for the dynamic management of free-floating bike-sharing systems. Transportation Research Part C: Emerging Technologies 87, 159–182.

Caggiani, L., Ottomanelli, M., 2013. A dynamic simulation based model for optimal fleet repositioning in bike-sharing systems. Procedia-Social and Behavioral Sciences 87, 203–210.

Chemla, D., Meunier, F., Calvo, R.W., 2013. Bike sharing systems: Solving the static rebalancing problem. Discrete Optimization 10, 120–146.

DeMaio, P., 2009. Bike-sharing: History, impacts, models of provision, and future. Journal of public transportation 12, 41–56.

Ghamami, M., Shojaei, M., 2018. Introducing a design framework for a multi-modal public transportation system, focusing on mixed-fleet bike-sharing systems. Transportation Research Record 2672, 103–115.

Jiménez-Meroño, E., Soriguera, F., 2024. Optimization of bike-sharing repositioning operations: a reactive real-time approach. EURO Journal on Transportation and Logistics 13, 100138.

Kloimüllner, C., Papazek, P., Hu, B., Raidl, G.R., 2014. Balancing bicycle sharing systems: an approach for the dynamic case, in: European conference on evolutionary computation in combinatorial optimization, Springer. pp. 73–84.

Legros, B., 2019. Dynamic repositioning strategy in a bike-sharing system; how to prioritize and how to rebalance a bike station. European Journal of Operational Research 272, 740–753.

Li, Y., Szeto, W.Y., Long, J., Shui, C.S., 2016. A multiple type bike repositioning problem. Transportation Research Part B: Methodological 90, 263–278.

Martins Silva, M.C., Aloise, D., Jena, S.D., 2023. Towards effective rebalancing of bike-sharing systems with regular and electric bikes .

O'Brien, O., 2025. Bike sharing world map. URL: `https://bikesharingworldmap.com`.

O'Mahony, E., Shmoys, D., 2015. Data analysis and optimization for (citi) bike sharing, in: Proceedings of the AAAI conference on artificial intelligence.

Pfrommer, J., Warrington, J., Schildbach, G., Morari, M., 2014. Dynamic vehicle redistribution and online price incentives in shared mobility systems. IEEE Transactions on Intelligent Transportation Systems 15, 1567–1578.

Raviv, T., Tzur, M., Forma, I.A., 2013. Static repositioning in a bike-sharing system: models and solution approaches. EURO Journal on Transportation and Logistics 2, 187–229.

Regue, R., Recker, W., 2014. Proactive vehicle routing with inferred demand to solve the bikesharing rebalancing problem. Transportation Research Part E: Logistics and Transportation Review 72, 192–209.

Shui, C.S., Szeto, W.Y., 2018. Dynamic green bike repositioning problem–a hybrid rolling horizon artificial bee colony algorithm approach. Transportation Research Part D: Transport and Environment 60, 119–136.

Shui, C.S., Szeto, W.Y., 2020. A review of bicycle-sharing service planning problems. Transportation Research Part C: Emerging Technologies 117, 102648.

Soriguera, F., Jiménez-Meroño, E., 2020. A continuous approximation model for the optimal design of public bike-sharing systems. Sustainable Cities and Society 52, 101826.

Wang, Y., Szeto, W.Y., 2018. Static green repositioning in bike sharing systems with broken bikes. Transportation Research Part D: Transport and Environment 65, 438–457.

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

Zhu, S., 2021. Optimal fleet deployment strategy: Model the effect of shared e-bikes on bike-sharing system. Journal of Advanced Transportation 2021, 6678637.