A Multi-Stage Optimization Framework for Battery Swapping in Urban E-Bike Sharing Systems

Fatemeh Farshchi Heydari

A Thesis Submitted to the Faculty of Graduate Studies

In Partial Fulfillment of the Requirements

for the Degree of Master of Applied Science

Graduate Program in Civil Engineering

York University

Toronto, Ontario

January 2026

Abstract

Shared electric bicycles (e-bikes) are increasingly central to urban transportation, providing a faster, more accessible alternative to conventional bicycles for short trips. As their use grows, maintaining sufficient battery levels across large fleets becomes critical. Among battery management strategies, battery swapping, which replaces depleted batteries with fully charged ones, offers a scalable solution, especially for dockless systems that lack fixed charging stations. However, it presents logistical challenges related to vehicle routing, battery delivery, and capacity constraints. This thesis introduces an optimization framework that combines dynamic clustering with a dual-objective vehicle routing model. E-bikes requiring service are grouped into van-feasible clusters based on battery level, spatial proximity, and fleet availability. Each cluster is then optimized to minimize travel distance while maximizing economic return. The framework is applied to real-world data from San Francisco's Bay Wheels system, accessed via the General Bikeshare Feed Specification (GBFS). Results show that the proposed method reduces travel distance and enhances the efficiency of battery-swapping operations.

Dedication

To my family, whose unwavering love and support made this journey possible.

And to the city streets, bikes, and data, for inspiring every line of this work.

Acknowledgements

I would like to express my deepest gratitude to my supervisor, Dr. Mehdi Nourinejad, for their guidance, insight, and encouragement throughout this research. I am also thankful to Interactive-OR lab for providing the necessary resources and academic environment that supported my work. Special thanks go to my colleague and friend Dr. Ghazaleh Mohseni who offered their feedback, support, and camaraderie throughout the process. I am especially grateful to my family for their constant encouragement, patience, and belief in me during this journey.

Contents

List of Tables

List of Figures

1 Introduction

Shared electric bicycles (e-bikes) are rapidly redefining short-distance urban travel by extending the reach and appeal of micromobility systems. Their ability to overcome physical barriers such as hills, longer distances, and rider fatigue has made sustainable mobility accessible to a broader segment of the population, for a wider range of trip purposes [2]. In 2023, shared micromobility systems across the United States and Canada recorded 157 million trips, a 20% increase over the previous year, with e-bikes driving much of this growth [1]. In the United States, station-based e-bike trips increased by 40% (from 20 million in 2022 to 28 million), accounting for 46% of all station-based bike share trips. Dockless e-bikes experienced similarly strong growth, with trips rising nearly 50% (from 4.5 million to 6.7 million), and Canadian dockless e-bike usage nearly quadrupling to 300,000 trips. These trends underscore the increasing role of e-bikes in shaping the future of urban mobility. However, this rapid expansion introduces new layers of operational complexity, requiring adaptive and efficient strategies to support system reliability and scalability.
Globally, cities are increasingly investing in micromobility as a sustainable and space-efficient alternative to private car travel. E-bike sharing systems play a key role in this shift, offering low-emission, flexible transportation options that complement public transit and reduce urban congestion. However, the scalability of these systems depends on operational strategies that are both cost-effective and environmentally responsible. Battery management is one of the most critical bottlenecks in this context, making it an essential area of research for cities pursuing climate and mobility goals.
Battery availability is a critical determinant of service quality in shared e-bike systems. Low state-of-charge (SOC) levels reduce trip flexibility, contribute to range anxiety, and increase user dissatisfaction [3]. These user-level concerns have broader implications for system performance: an undercharged fleet can result in decreased ridership, increased service complaints, and diminished reliability. From the operator's perspective, maintaining battery availability requires continuous monitoring, collection, and redistribution of vehicles. These tasks are labor-intensive, time-sensitive, and costly. As a result, battery management plays a central role in ensuring the efficiency, equity,

and long-term sustainability of shared e-bike services.

A variety of operational strategies have been implemented to manage battery levels in shared systems, each with unique infrastructure, labor, and cost implications. One common approach involves equipping docking stations with charging capabilities, allowing idle e-bikes to recharge while parked. This strategy aligns with natural usage patterns but demands significant capital investment, including high-capacity electrical connections and access to the power grid. In dense urban environments, such infrastructure can be costly to install and difficult to maintain. An alternative strategy is to collect low-SOC e-bikes and transport them to centralized depots for charging and reallocation. While this reduces infrastructure needs, it imposes high operational costs due to labor, travel time, and delays in bike availability. More recently, battery swapping has emerged as a promising solution, particularly for dockless systems that lack fixed charging stations. In this approach, depleted batteries are removed and replaced with fully charged ones on-site. This method minimizes downtime and increases fleet availability by separating the charging process from the vehicle. However, battery swapping introduces its own set of logistical challenges, including inventory control, real-time routing, and workforce deployment.

Existing research has explored various dimensions of battery management in e-bike systems. Studies on fixed charging infrastructure have largely focused on optimizing station placement and grid connectivity to meet spatial demand. For instance, Bassolas et al. [4] analyzed the spatiotemporal dynamics of SOC patterns in Barcelona's bike-sharing system and proposed predictive models that leverage mobility data to forecast battery depletion. Meanwhile, research on battery swapping has advanced both the technological and operational aspects of implementation. Wu [5] categorized battery swapping systems by decision-making level, while Shao et al. [6] introduced a three-stage optimization framework involving clustering, routing, and vehicle assignment. Additional work has used reinforcement learning to optimize real-time swapping and rebalancing decisions in large networks [7], queuing models to reduce user delays [8], and stochastic Markov models to anticipate energy depletion events [9]. However, the majority of these studies focus on structured, station-based networks or rely on static optimization models. Few address the dynamic operational

demands of dockless systems, where spatial distribution is more diffuse, SOC variability is higher, and system responsiveness is more critical.

This thesis addresses the problem of battery swapping in large-scale, dockless e-bike sharing systems. Specifically, it proposes an optimization framework to coordinate the delivery of fully charged batteries from a central depot to depleted bikes while minimizing operational costs and maximizing economic return. The problem is formulated as a variant of the vehicle routing problem (VRP), where fully charged batteries represent the goods, and the delivery vehicle is a battery-swapping van. To manage complexity and improve scalability, the system incorporates a dynamic clustering phase that partitions e-bikes into feasible service clusters. Clustering is performed based on factors such as SOC levels, geographic proximity, the number of available vans, and their carrying capacity. Each resulting cluster is designed so that a single van can service all bikes within it without exceeding operational constraints. A dual-objective optimization model is then applied to each cluster, aiming to reduce total travel distance while maximizing profitability.

The objective of this research is to develop an integrated optimization framework for battery swapping operations in dockless e-bike sharing systems that explicitly combines spatial clustering with vehicle routing optimization. The study aims to use clustering as a mechanism to reduce problem scale and partition large, spatially diffuse fleets into operationally feasible service regions, within which routing decisions can be solved efficiently. Within each cluster, the battery swapping operation is formulated as a vehicle routing problem that accounts for battery state-of-charge, expected future demand, and vehicle capacity constraints. By linking clustering and routing into a unified workflow, the proposed approach seeks to improve computational tractability while balancing operational efficiency and economic performance in large-scale urban systems.

San Francisco was selected as the primary case study for this research due to the availability of high-quality, real-time data from its Bay Wheels bike-sharing system, as well as its adoption of a dockless operational model. The city's dense, topographically varied urban environment, coupled with its emphasis on sustainable transportation, makes it a representative and practically relevant setting for evaluating the proposed battery swapping framework.

The remainder of this thesis is organized as follows. Chapter 2 reviews the existing literature on battery management strategies in shared e-bike systems. Chapter 3 presents the problem formulation and profitability model. Chapter 4 describes the solution methodology, including the clustering algorithm and routing approach. Chapter 5 presents the experimental results. Chapter 6 details the San Francisco case study. Finally, Chapter 8 concludes the thesis with key findings and suggestions for future work.

Table 1: Comparison of battery management strategies for shared e-bike systems

Strategy	Advantages	Limitations
Station-Based Charging	• Charging occurs passively while bikes are docked. • Leverages existing user behavior (returning to stations). • Reduces need for frequent operator interventions.	• Requires significant infrastructure investment. • Limited to station-based systems; not applicable to dockless fleets. • Space and permitting constraints in dense urban cores.
Depot Charging	• Centralized control of charging operations. • Lower reliance on distributed infrastructure. • Simplifies monitoring and maintenance at a single location.	• High labor and vehicle transport costs. • Delay between retrieval and redeployment reduces fleet availability. • Inefficient for large, spatially distributed systems.
Battery Swapping	• Enables immediate redeployment of e-bikes. • Decouples charging from vehicle use, improving availability. • Flexible and mobile; particularly suited for dockless fleets.	• Requires careful inventory management of charged/depleted batteries. • Complex routing and scheduling of service vans. • Labor and coordination challenges increase with fleet size.

2 Background and Literature Review

This chapter provides the conceptual and empirical foundations for the thesis by combining background context with a focused review of the relevant literature. The discussion begins with the evolution and global diffusion of micromobility, highlighting the role of shared e-bikes, their integration with public transit, and the impacts of technology, governance, and equity on system performance. It then turns to operational and methodological perspectives, reviewing strategies for battery management, vehicle routing, and clustering in large-scale urban networks. Bringing together insights from transportation planning and operations research allows for a comprehensive understanding of both the societal role of shared micromobility and the technical challenges of sustaining it. The chapter concludes by identifying research gaps that motivate the optimization framework developed in subsequent chapters.

2.1 Micromobility and the Role of E-Bikes

Micromobility refers to the set of lightweight, often electrically assisted transport modes designed for short-distance urban trips. Shared micromobility systems, which include bicycles, e-bikes, and scooters, have expanded rapidly in the past decade due to their potential to reduce congestion, cut emissions, and provide flexible last-mile connectivity [1]. Among these, electric bicycles (e-bikes) have emerged as a particularly important mode because they extend the feasible trip range and reduce physical barriers such as steep gradients and rider fatigue [2].

Recent adoption statistics demonstrate the transformative effect of e-bikes on travel behaviour. In the United States and Canada, shared micromobility reached 157 million trips in 2023, with e-bikes accounting for a growing proportion of total usage. Dockless systems, in particular, recorded nearly 50% year-over-year growth, underscoring their scalability and popularity in dense urban centers. These figures indicate that e-bikes are no longer a supplementary option but rather a core component of emerging urban mobility ecosystems(Figure 2).

The concept of bikesharing has undergone a remarkable evolution over the past six decades, pro-

gressing through distinct technological and organizational "generations." Each stage reflects broader shifts in urban planning, transport policy, and digital technology, moving from grassroots social experiments to globally integrated mobility systems. A detailed understanding of this historical trajectory provides valuable context for contemporary e-bike fleets and their operational challenges.

First Generation: Free Bikes The first generation of bikesharing began in Amsterdam in 1965 with the launch of the "White Bike Plan" (*Witte Fietsen*). Fifty bicycles were painted white and distributed freely across the city for public use. The idea, rooted in countercultural activism, was to provide an accessible, environmentally friendly form of collective mobility. However, the absence of tracking or enforcement mechanisms led to rapid system failure: many bicycles were stolen, vandalized, or appropriated for private use, and the program collapsed within days [10]. Despite its short-lived nature, this initiative established the symbolic and practical foundation for bikesharing as a novel mode of urban transport.

Second Generation: Coin-Deposit Systems A second generation emerged in the early 1990s, beginning with pilot schemes in Denmark (Farso, Grenra, and Nakskov) and scaling up with the launch of Copenhagen's Bycyklen (City Bikes) program in 1995. These systems introduced specialized bicycles designed for heavy-duty use, docking infrastructure, and a refundable coin deposit to reduce theft. Bycyklen quickly became a landmark case, offering hundreds of bicycles across central Copenhagen. While the use of coin deposits represented a major improvement over the free-bike model, anonymity of users persisted and theft remained problematic [10, 11]. Nevertheless, the second generation demonstrated that bikesharing could operate as a formalized, city-supported service rather than a grassroots experiment.

Third Generation: ICT-Enabled Bikesharing The third generation of bikesharing marked a decisive transformation by integrating information and communication technologies (ICT). Early examples included Portsmouth University's Bikeabout program in 1996, which used magnetic stripe cards, and Rennes' Vélo à la Carte in 1998, which introduced smartcards and electronic locking systems. The turning point came with Lyon's Velo'v in 2005, operated by JCDecaux, which deployed 1,500 bicycles and quickly attracted 15,000 members with usage rates averaging 6.5 trips

per day. Paris soon followed with the launch of Vélib' in 2007, offering 7,000 bicycles that rapidly expanded to more than 20,000. Vélib's visibility and popularity captured international attention, positioning bikesharing as a mainstream urban transport mode rather than a niche experiment [10]. Third-generation systems introduced several innovations that remain central to bikesharing today: smartcards and fobs for user identification, GPS-enabled tracking, automated docking stations, and integration with municipal transport authorities. By 2009, more than 120 ICT-based bikesharing programs were operating globally, including North America's first large-scale deployments such as BIXI in Montreal and Capital Bikeshare in Washington, D.C. [11]. The use of ICT reduced theft, enabled detailed ridership data collection, and made large-scale operations financially viable.

Fourth Generation: Integration, Electrification, and Flexibility Building on the success of ICT-enabled systems, scholars and practitioners identified the emergence of a fourth generation of bikesharing in the 2010s [12]. This generation is characterized by modular and mobile docking infrastructure, solar-powered stations, integration with public transport smartcards and mobile applications, and increasingly, the inclusion of electric bicycles (pedelecs). Operational innovations, such as incentive-based redistribution, where users are rewarded for returning bicycles to underserved stations, have emerged to reduce the costly manual rebalancing. Fourth-generation systems are also defined by new business models, including public–private partnerships, advertising-based financing, and for-profit dockless operators.

The rise of dockless bikesharing, pioneered in China by companies such as Ofo and Mobike in the mid-2010s, further extended fourth-generation principles by removing the need for fixed stations altogether. Users could unlock bicycles with smartphone apps and leave them at virtually any location within a geofenced service area. While dockless systems scaled rapidly, reaching millions of bicycles in Chinese cities within a few years, they also introduced new challenges, including oversupply, sidewalk clutter, and financial instability. These experiences underscored the importance of governance frameworks and regulatory oversight in shaping sustainable system growth [13].

Impacts and Legacy Across its historical evolution, bikesharing has demonstrated significant societal benefits. Studies document increases in cycling mode share in cities such as Barcelona and

Paris following the launch of third-generation programs, improvements in public transport connectivity through first–last mile trips, and reductions in greenhouse gas emissions from car substitution [10]. The successive generations illustrate a pattern of technological adaptation and institutional learning: each new wave of systems addressed the shortcomings of its predecessors while expanding bikesharing's scale, scope, and integration into the wider transport system.

From the early free-bike experiments of the 1960s to today's integrated, ICT-enabled, and increasingly electrified fleets, bikesharing has become an essential element of urban mobility. This historical trajectory highlights how micromobility has transitioned from peripheral experiments to integral components of transport ecosystems. Understanding this evolution not only contextualizes the rise of shared e-bikes but also underscores why challenges such as battery management and routing are central to the sustainability of future generations of micromobility systems.

The spread of bikesharing has been profoundly shaped by regional transport cultures, governance traditions, and technological contexts. While the underlying concept of providing bicycles for shared public use is universal, the scale, organization, and impacts of bikesharing differ considerably across Europe, North America, and Asia. Understanding these regional contrasts is essential for situating the contemporary rise of shared e-bikes and the operational challenges they pose.

Europe has played a leading role in embedding bikesharing into mainstream transport policy. Long before the ICT-enabled systems of the 2000s, many European cities were already at the forefront of promoting cycling as a daily mode of travel. Policy frameworks in countries such as the Netherlands, Denmark, and Germany treated cycling not as a recreational activity but as a core part of the transport system, supported by dense networks of cycle tracks, integrated land use planning, and financial incentives for sustainable modes [14]. Against this background, bikesharing was introduced as a complement to existing cycling culture, supported by pro-cycling governments and public-private partnerships.

European cases illustrate the ability of policy design to shape bikesharing outcomes. Programs such as Lyon's Velo'v (2005) and Paris' Vélib' (2007) achieved rapid uptake partly because they were embedded in a broader political strategy to reduce congestion, reclaim street space, and integrate

with public transport [15]. Mode shift impacts were significant: in Paris, cycling mode share increased from 1 percent in 2001 to 2.5 percent by 2007 after the launch of Vélib', while in Barcelona, the introduction of Bicing raised cycling's share from 0.75 percent in 2005 to 1.76 percent two years later [10]. European systems also demonstrated the importance of governance models: many were operated through advertising concessions (e.g., JCDecaux in Paris) or quasi-governmental transport agencies, ensuring long-term integration into transport networks. The European experience highlights that bikesharing thrives when supported by strong cycling infrastructure, multi-modal integration, and explicit government backing.

In contrast, bikesharing in North America emerged later, with Montreal's BIXI in 2009 marking the first large-scale third-generation system on the continent. Subsequent programs such as Capital Bikeshare (Washington, D.C.) and Citi Bike (New York City) expanded quickly, demonstrating demand for cycling even in cities with historically low cycling mode share. According to Fishman's comprehensive review of the bikesharing literature, the rapid growth of these systems underscores their potential to reduce reliance on automobiles, provide first–last mile connections, and generate environmental benefits [15].

Fishman identifies several consistent global impacts of bikesharing. First, bikesharing increases access to cycling among new users, broadening the socio-demographic base of urban cycling. Second, systems consistently generate positive environmental outcomes, with reductions in greenhouse gas emissions and air pollutants when bikeshare trips substitute for car trips. Third, integration with public transport is a recurring theme: bikesharing often substitutes for short transit trips while simultaneously extending the reach of rail and bus networks by providing first–last mile connections. However, Fishman also emphasizes that challenges remain, particularly regarding financial sustainability, equity of access, and the environmental footprint of redistribution activities. This synthesis positions bikesharing as a globally significant innovation while highlighting the need for regionally adapted operational models. Perhaps the most dramatic regional development in bikesharing has been the rapid rise of dockless systems in China. Beginning in the mid-2010s, companies such as Mobike and Ofo deployed millions of bicycles across major Chinese cities, bypassing the fixed

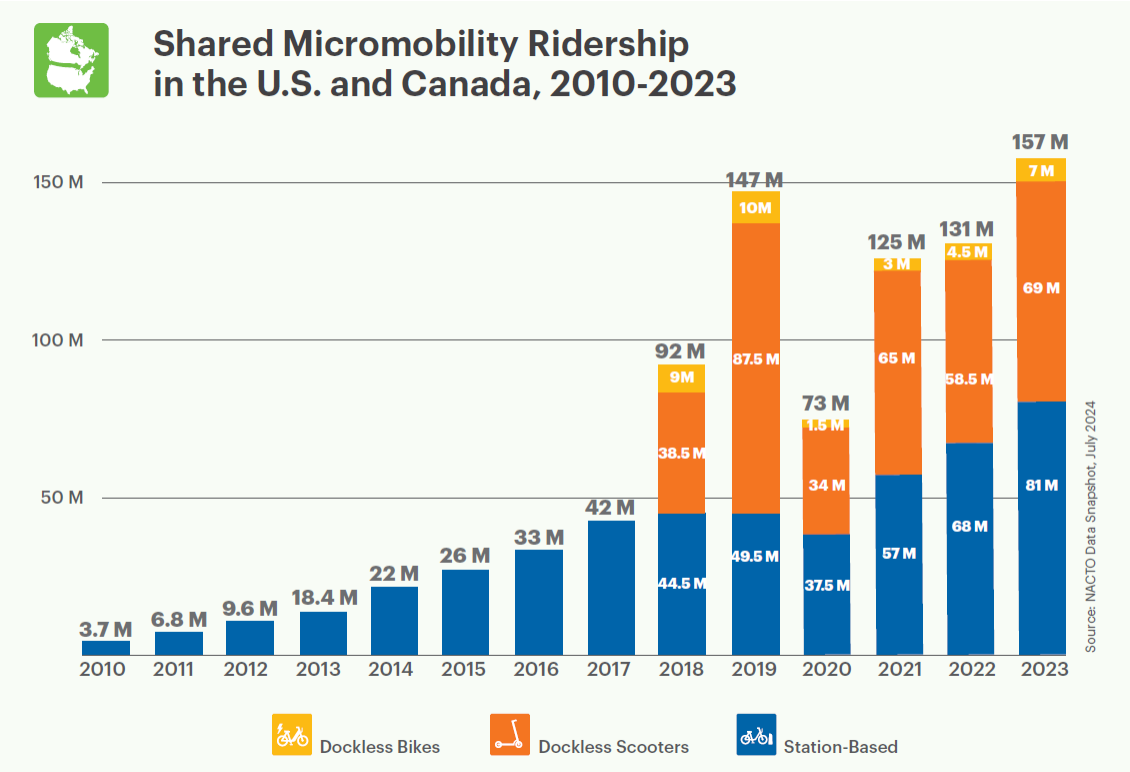

Figure 1: Shared micromobility ridership in the U.S. and Canada has grown nearly fortyfold since 2010, with e-bikes and scooters now accounting for a substantial share alongside traditional station-based bikes. Source: Nacto Data Snapshot [1]

docking infrastructure that characterized earlier systems. Instead, users could unlock bikes with smartphone applications and leave them anywhere within a service area, enabled by GPS tracking and mobile payments. This model, pioneered in China, represented a step change in convenience and scalability, transforming bikesharing from a niche service into a ubiquitous feature of the urban landscape [16].

Evidence from Beijing shows that dockless bikesharing addressed critical first–last mile gaps in the transport system. Surveys and interviews revealed that users were predominantly younger, educated, and technologically adept, and that dockless bikes provided time savings and improved connectivity to metro and bus systems. At the same time, the systems appealed to a diverse income base, lowering barriers to entry for cycling in congested urban environments [16]. Environmental impacts were significant: in Shanghai, dockless bikesharing was estimated to have saved thousands of tons of gasoline and tens of thousands of tons of CO_2 emissions within a single year. These

outcomes illustrate the potential of large-scale dockless deployment to contribute meaningfully to sustainability goals.

However, China's experience also revealed serious governance and operational challenges. Oversupply of bicycles led to images of "bike graveyards" as firms flooded the streets with more bikes than cities could accommodate. Problems of vandalism, abandoned or broken bikes, and unsafe riding behavior strained public tolerance. Furthermore, the financial sustainability of private operators proved precarious, with both Mobike and Ofo experiencing severe losses and withdrawals from many markets. Sun's study highlights the need for multi-party governance involving city authorities, private firms, and the public to ensure that dockless bikesharing can be sustained as a long-term mobility option [16].

Taken together, these regional experiences illustrate the varied trajectories of bikesharing across the globe. European systems benefited from supportive policy environments and long-standing cycling cultures, embedding bikesharing within integrated transport strategies. North American systems demonstrated the global diffusion of ICT-enabled bikesharing and provided systematic evidence of its benefits and challenges across diverse urban contexts. China's dockless boom showcased the potential for rapid scale-up and innovation but also highlighted the risks of market-driven oversupply and insufficient regulation. These comparative insights reinforce the argument that bikesharing is not a uniform technology but a flexible transport mode whose success depends on local institutional, cultural, and governance conditions.

For the purposes of this thesis, these global trends underscore the importance of viewing shared e-bikes not as isolated systems but as part of a broader international evolution. The lessons from Europe, North America, and China demonstrate both the opportunities and challenges of integrating bikesharing into sustainable urban mobility. They also highlight why operational issues such as battery management and routing, while technical in nature, are inseparable from the wider social, environmental, and governance contexts in which shared micromobility systems operate.

One of the most widely recognized benefits of shared micromobility is its potential to enhance public transit by solving the long-standing "first–last mile" problem. Public transport networks are

highly efficient along trunk corridors such as metro, bus rapid transit, and commuter rail lines, but their accessibility is often limited by the need for connecting trips at either end. Bikesharing offers a flexible, low-cost, and spatially efficient solution to bridge these gaps, reducing access and egress times and expanding the effective catchment area of transit stations.

Research consistently highlights this complementary relationship. Shaheen and Chan provide a broad policy and conceptual overview of how shared mobility, including bikesharing, carsharing, ridesourcing, and microtransit, facilitates first–last mile connections to public transit [17]. They argue that integration is increasingly enabled by advances in information and communication technology (ICT), such as smartphone applications and multimodal payment systems, which collectively lay the foundation for Mobility-as-a-Service (MaaS). Within this framework, bikesharing is not a stand-alone system but part of a larger ecosystem of shared mobility options designed to reduce automobile dependence and enhance multi-modal travel flexibility.

Empirical studies reinforce this integrative role. Ma et al. model the interactions between bikesharing and public transit ridership, finding that bikesharing acts as a substitute for short bus trips while simultaneously complementing longer metro and rail journeys [18]. By providing efficient access to stations in areas with weak feeder bus service, bikesharing strengthens the overall performance of the transit network. Importantly, the study shows that the net effect is positive: bikesharing extends the reach of transit systems, enabling users to complete multimodal journeys more quickly and reliably. This evidence highlights that bikesharing is not merely recreational, but structurally tied to public transit demand.

Case studies from Europe confirm these findings. Jäppinen et al. use accessibility modeling in Greater Helsinki to estimate the potential impact of bikesharing on public transport travel times [19]. They report average travel time savings of approximately 10 percent (equivalent to six minutes per trip) when bikesharing was introduced as a feeder mode. The benefits were greatest in suburban areas where conventional feeder buses were infrequent, demonstrating that shared bicycles can significantly improve the attractiveness of rail-based commuting. This work illustrates the importance of contextualizing bikesharing not only as a cycling policy but also as a strategic

extension of high-capacity public transport.

Taken together, these studies demonstrate that bikesharing systems consistently enhance the accessibility, efficiency, and appeal of public transit. At the same time, they underscore the operational dependency of this integration on reliable fleet availability. If e-bikes are unavailable or in a depleted battery state, their ability to serve as transit feeders is compromised, reducing both system efficiency and user satisfaction. For this reason, the operational challenge of battery management, addressed through strategies such as swapping and routing optimization, is not only a technical issue but also a central determinant of whether bikesharing can realize its role within MaaS and the wider sustainable mobility ecosystem.

The past decade has seen rapid global adoption of shared micromobility systems. By 2019, more than 1,600 programs operated worldwide, covering both large metropolitan areas and smaller cities [20]. In North America, ridership has grown particularly fast: since the launch of Montréal's BIXI in 2009, shared micromobility users have taken more than half a billion trips. Despite the disruption of the COVID-19 pandemic, ridership rebounded strongly, with 112 million trips in 2021, nearly returning to pre-pandemic levels [21], as can be seen in Figure 1. Importantly, e-bikes now account for a rapidly increasing share of trips, reflecting their ability to expand the effective catchment area of bike sharing. Systematic reviews confirm that shared micromobility has emerged as one of the most visible and transformative urban transport innovations in the 21st century [22].

The COVID-19 pandemic provided a natural stress test for urban transport systems, highlighting both vulnerabilities and adaptive capacities. Public transit systems worldwide experienced dramatic ridership declines due to social distancing mandates, remote work policies, and concerns over contagion in enclosed spaces. In contrast, shared micromobility demonstrated relative resilience, with many cities observing that bikesharing and e-scooter services were able to maintain a significant portion of their demand, adapt quickly to changing conditions, and even expand their role as essential transport services.

Evidence from New York City illustrates this resilience vividly. Teixeira and Lopes analyzed ridership patterns in Citi Bike, the city's bikesharing system, compared to the subway during the first

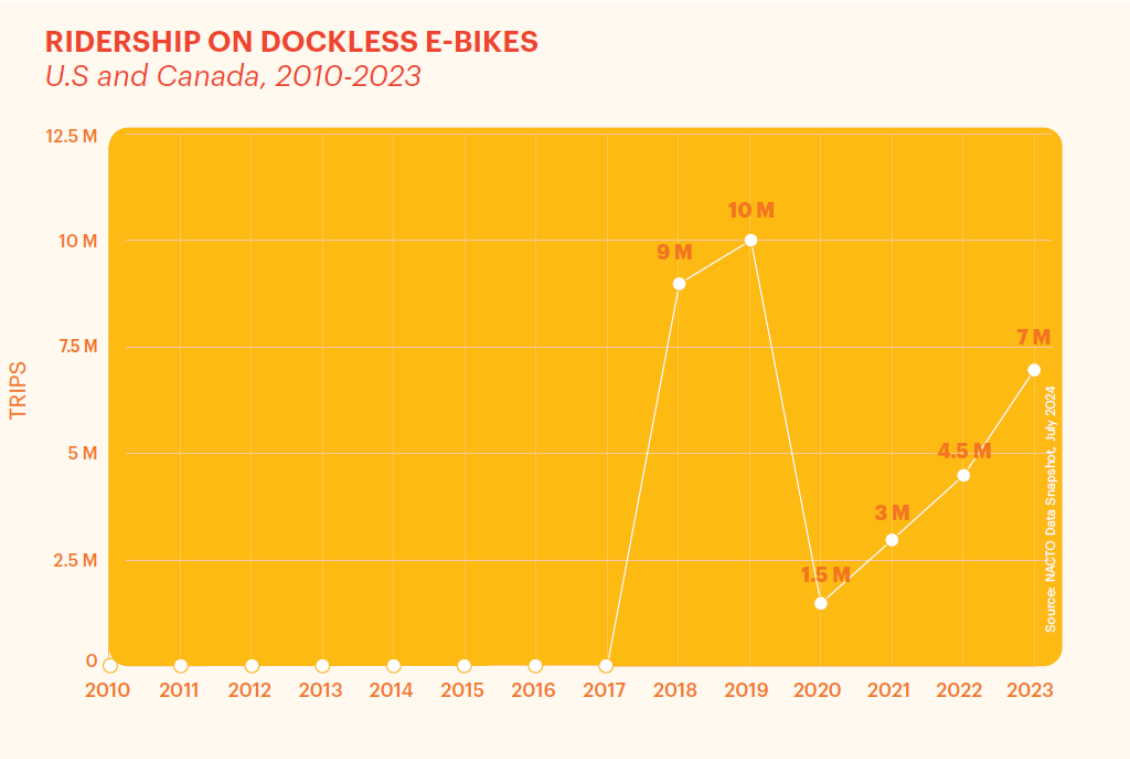

Figure 2: Dockless e-bike ridership in the U.S. and Canada has grown significantly since 2017, demonstrating rapid adoption, temporary pandemic decline, and post-2020 recovery. Source: Nacto Data Snapshot [1]

months of the pandemic [23]. While subway ridership fell by more than 90 percent in March 2020, Citi Bike demand declined by a smaller margin of approximately 71 percent. Moreover, average trip durations increased by more than 67 percent, indicating that users were undertaking longer, essential trips rather than short leisure rides. These patterns suggest that bikesharing partially substituted for disrupted subway travel, particularly for essential workers seeking socially distanced alternatives. The findings underscore that bikesharing systems can act as a flexible complement to traditional transit, even under severe external shocks.

Broader cross-country evidence confirms that the impacts of the pandemic on bikesharing were heterogeneous and shaped by local policy responses, cultural attitudes, and system design. Nikiforiadis et al. compared bikesharing demand across several European cities and found that while ridership initially declined sharply during lockdown periods, recovery trajectories varied considerably [24]. Cities with robust cycling infrastructure, supportive local policies, and integrated micromobility ecosystems rebounded more quickly, while others struggled with sustained declines.

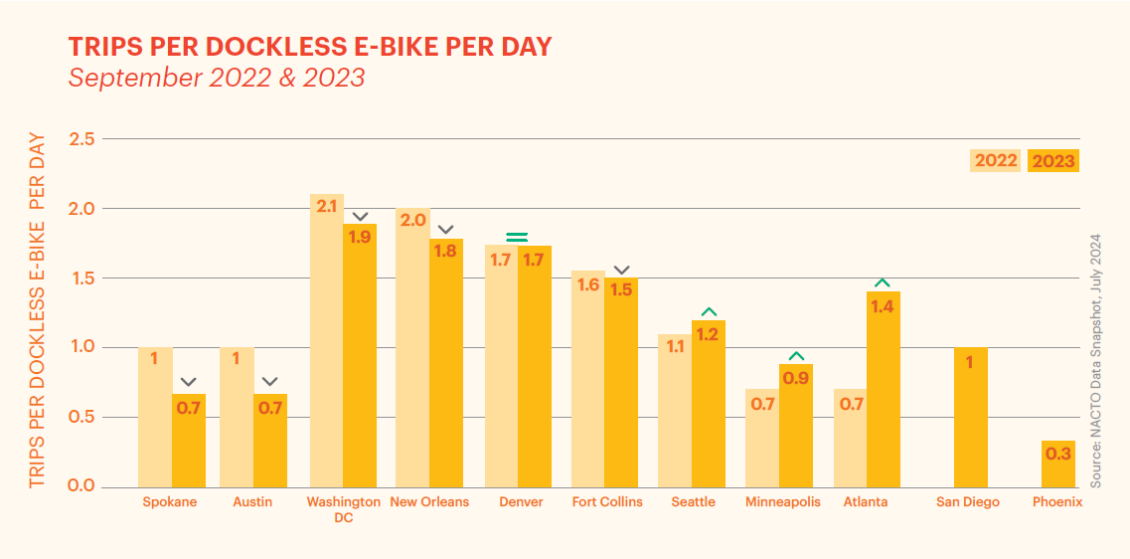

Figure 3: Average daily trips per dockless e-bike vary widely across cities, reflecting local travel demand, cycling infrastructure, and system design. Source: Nacto Data Snapshot [1]

The study highlights that bikesharing's resilience is not automatic but contingent upon governance structures and contextual factors that determine its ability to adapt during (Figure 3).

In North America, bikesharing's role as an "essential service" became especially evident. According to the National Association of City Transportation Officials (NACTO), shared micromobility systems in the United States recorded 65 million trips in 2020 despite widespread restrictions, with ridership rebounding to 112 million trips in 2021, nearly reaching pre-pandemic levels [21]. Station-based bikeshare proved more resilient than dockless scooters, reflecting both infrastructure stability and integration with transit networks. Operators introduced emergency measures such as reduced-cost passes for frontline workers, expanded stations near hospitals, and enhanced cleaning protocols. Importantly, the pandemic also catalyzed long-term behavioural changes: ridership shifted away from traditional commuting peaks toward more evenly distributed patterns throughout the day, and the share of casual users increased. The surge in e-bike ridership—growing nearly thirty-fold between 2018 and 2021 was particularly notable, as can be seen on Figures 1 and 2, reinforcing the importance of electrification for resilience and recovery.

Taken together, these studies demonstrate that bikesharing systems displayed remarkable adaptability during the pandemic. While not immune to demand shocks, bikesharing was less severely

affected than traditional transit and often stepped in as a crucial alternative for essential mobility. The pandemic also accelerated pre-existing trends, including the mainstreaming of e-bikes, the diversification of user bases, and the recognition of micromobility as a critical component of resilient urban transport. For this thesis, these insights emphasize that effective operational strategies such as battery swapping and routing optimization are not only technical necessities but also resilience measures, ensuring that e-bike fleets remain available to support urban mobility under both normal and crisis conditions.

Beyond operational and behavioral dimensions, the environmental implications of e-bikes and shared micromobility have been central to academic and policy debates. Early studies in China framed e-bikes as both an opportunity and a challenge. Cherry and Cervero [25] showed that e-bikes largely substitute for motorized trips, particularly buses and private cars, thereby reducing per-capita emissions and congestion. However, they also noted that the net benefits depend on the electricity mix and on responsible management of lead-acid batteries, which at the time dominated the Chinese market. This dual framing, environmental opportunity versus risk, has shaped subsequent evaluations of e-bike sustainability worldwide.

More recent meta-analyses provide systematic evidence of substitution effects across contexts. Bigazzi and Wong [26] synthesized 24 empirical studies and found that e-bikes most frequently displace public transit trips (33%), followed by conventional bicycles (27%), automobiles (24%), and walking (10%). Importantly, regional variation emerged: in China, e-bikes tend to replace bus travel, while in North America and Europe, they more often displace private car use. These differences underscore how local travel cultures and infrastructure mediate the environmental outcomes of e-bike adoption. Encouragingly, newer studies reveal a trend toward greater substitution of automobiles and walking and less displacement of cycling, suggesting that as e-bikes mainstream, their marginal environmental benefits improve.

From a system-level perspective, life-cycle assessments (LCAs) highlight that the sustainability of bikesharing depends not only on user substitution patterns but also on system operations. Luo et al. [27] compared docked and dockless bikesharing systems and found that manufacturing impacts, re-

distribution logistics, and battery management dominate life-cycle emissions. Dockless fleets often entail higher operational costs and environmental burdens due to frequent rebalancing, but targeted interventions, such as battery swapping strategies, can mitigate these impacts. These findings reinforce the importance of operational efficiency, including routing and clustering optimization, as a critical determinant of environmental performance.

Together, these studies establish that e-bikes offer substantial energy and environmental advantages when they displace motorized modes, but that outcomes depend strongly on contextual factors such as the local electricity grid, fleet management practices, and system design. For this thesis, the environmental framing motivates battery swapping optimization not only as an operational necessity but also as a pathway to reduce the life-cycle footprint of shared e-bike systems.

Electrification is transforming the scope and function of micromobility. Shared e-bikes reduce physical barriers such as steep gradients and rider fatigue, thereby broadening the demographic base of potential users. Evidence from San Francisco shows that dockless e-bikes are used for longer trips and more challenging topographies compared to station-based pedal bikes [28]. Nationwide, e-bike ridership nearly tripled between 2018 and 2021, with more than 14.5 million e-bike trips recorded in 2021 [21]. This trend has profound operational implications: batteries must be charged, swapped, and redistributed efficiently to maintain availability. As more fleets transition to fully electric models, battery management becomes not only a logistical challenge but also the central determinant of service quality.

The impacts of shared micromobility are multifaceted. Environmentally, studies show substantial reductions in fuel use and carbon emissions when bike-share trips substitute for car travel; in Shanghai, bike sharing saved more than 25,000 tons of CO_2 in a single year [29]. Life-cycle assessments also highlight that while dockless systems generate higher rebalancing emissions, improved logistics and vehicle lifespans can offset these impacts [27]. Safety outcomes are similarly promising: crash severity for bike-share users is lower than for private cyclists, suggesting a "safety in numbers" effect [30]. At the same time, equity concerns persist. Dockless pilots in Seattle showed relatively wide access, yet disadvantaged neighborhoods received fewer bikes and less frequent

rebalancing [31]. Broader critiques argue that mainstream systems often underserve marginalized groups, raising questions about inclusivity [32]. Finally, the built environment plays a critical role. Land-use mix, density, and street design strongly condition the uptake of shared bikes and scooters, emphasizing the need to integrate micromobility into planning at the node, link, and network levels [33].

Beyond environmental and operational benefits, bikesharing systems contribute significantly to public health by increasing population-level physical activity. Fishman et al. [34] conducted the first multi-city evaluation across Melbourne, Brisbane, Washington, D.C., London, and Minneapolis–St. Paul, finding that bikeshare consistently increases active travel minutes, particularly when replacing car or public transport trips, with overall positive net effects despite some substitution from walking.

Large-scale national evidence reinforces these findings. Fishman et al. [35] estimate that cycling in the Netherlands prevents approximately 6,500 deaths annually, extending life expectancy by half a year and generating economic health benefits equivalent to over 3% of GDP. These outcomes underscore the magnitude of health gains that can result when cycling becomes mainstream.

At the city level, Woodcock et al. [36] evaluated London's bikeshare program and showed that despite increased exposure to traffic risks and pollution, the physical activity gains produced substantial net health benefits, particularly among older age groups. Similarly, Otero et al. [37] highlight mortality reductions and public health savings associated with cycling across European cities, situating bikeshare as part of a broader health-promoting mobility policy.

Taken together, this body of research demonstrates that bikesharing is not merely a transport innovation but also a public health intervention. The capacity of bikeshare to induce additional minutes of active travel, reduce sedentary lifestyles, and extend life expectancy adds a critical dimension to the case for efficient and equitable management of e-bike fleets. Ensuring adequate vehicle availability and battery reliability therefore has implications not only for operational performance but also for urban health outcomes.

Table 2: Evolution of bikesharing systems across four generations

Generation	Key Features	Representative Examples	Limitations / Challenges
1st Generation (1960s–1970s)	Free community bicycles with no docking or tracking; grassroots and socially motivated schemes.	Amsterdam "White Bike Plan" (1965).	High theft and vandalism; lack of accountability led to rapid collapse.
2nd Generation (1990s)	Coin-deposit systems with basic docking and heavy-duty bicycles to deter theft.	Copenhagen Bycyklen (1995); Danish pilot schemes (Farso, Grenra, Nakskov).	Persistent theft due to anonymous use; limited scalability.
3rd Generation (late 1990s–2000s)	ICT-enabled systems with smartcards, GPS, and automated docks; user identification and data collection.	Bikeabout (Portsmouth, 1996); Vélo à la Carte (Rennes, 1998); Vélo'v (Lyon, 2005); Vélib' (Paris, 2007); BIXI (Montréal, 2009).	High infrastructure cost; reliance on docks limits flexibility.
4th Generation (2010s–present)	Mobile app integration, flexible or dockless operation, electrification, and incentive-based redistribution.	Dockless boom in China (Mobike, Ofo); U.S. dockless pilots (Seattle, 2017); electrified fleets (e.g., Bay Wheels, San Francisco).	Oversupply and financial instability; sidewalk clutter requiring regulation.

2.2 Policy and Equity Dimensions

The development of shared micromobility systems is shaped not only by technology and user demand, but also by the regulatory and institutional environments in which they operate. Governance determines the rules of access to public space, the responsibilities of operators, and the mechanisms for aligning private services with public goals. As such, policy frameworks provide the enabling conditions for bikesharing systems to expand, adapt, and sustain their role in urban transport. Before turning to operational challenges such as vehicle availability and battery logistics, it is therefore important to review the governance context that structures how shared mobility systems function in practice.

The expansion of shared micromobility has been accompanied by new governance frameworks. Cities employ regulatory levers such as device caps, service area boundaries, geo-fencing, and pricing structures to balance operator flexibility with public interests [13]. Many North American systems operate under public–private partnerships, which proved especially resilient during the pandemic [21]. Policy debates also extend to broader shared mobility ecosystems, where integration with public transport and Mobility-as-a-Service (MaaS) platforms is increasingly emphasized [17]. Governance therefore not only shapes system performance but also conditions how optimization frameworks, such as those developed in this thesis, can be implemented in practice.

Governance frameworks play a central role in shaping the development and sustainability of shared mobility systems. Evidence from carsharing provides a useful precedent for understanding the regulatory challenges and policy levers that now confront micromobility. Shaheen and Cohen [38] present one of the most comprehensive global reviews of carsharing, documenting its rapid expansion to more than 1,100 cities in 26 countries by 2010. Their analysis shows that supportive policies,such as access to on-street parking, streamlined insurance regulation, and partnerships with municipalities, were crucial in enabling growth. At the same time, governance frameworks had to address new challenges related to equity, environmental impacts, and competition with private automobile ownership.

These lessons are directly relevant to micromobility. The same policy questions, including the allocation of curb space, the management of the right of way, and the integration with public transport, now apply to shared bicycles and e-bikes. Furthermore, the entry of traditional rental car companies and automakers into the carsharing market illustrates how governance must adapt as private sector actors with significant resources enter shared mobility domains. In this way, the governance of carsharing can be seen as a precursor to the governance of bikesharing, offering both insights and cautionary lessons for policymakers tasked with regulating dockless and electrified micromobility systems.

While governance frameworks establish the institutional conditions under which shared mobility systems operate, their day-to-day performance ultimately depends on how operators manage

practical challenges such as vehicle availability, rebalancing, and battery logistics. These operational concerns are particularly acute in shared e-bike systems, where maintaining adequate state-of-charge levels across the fleet is central to ensuring reliability and user satisfaction. While governance frameworks establish the institutional rules and incentives that shape shared micromobility, the everyday success of these systems ultimately hinges on addressing practical operational challenges, most notably ensuring vehicle availability and managing battery logistics in shared e-bike fleets.

Equity has become a central concern in the expansion of shared micromobility systems, as studies consistently document disparities in both access and use. McNeil et al. [39] surveyed system operators, policymakers, and users across the United States and found that low-income and minority communities face persistent barriers to participation. These included cost, lack of payment flexibility, limited station coverage, and cultural perceptions of bikesharing. The study emphasized that overcoming these barriers requires targeted interventions such as subsidized memberships, cash-based payment options, and community engagement strategies to ensure that bikeshare systems serve a broader demographic.

Spatial analyses of accessibility provide further insight into equity outcomes. Desjardins et al. [40] evaluated Hamilton Bike Share's "Everyone Rides" initiative, which introduced twelve equity-focused stations in underserved neighborhoods. Using a balanced floating catchment area method, they found modest improvements in overall accessibility, but limited gains for the lowest-income households. The authors concluded that equity cannot be achieved simply by expanding system coverage; it also depends on the placement of stations, their capacity, and the integration of supportive cycling infrastructure.

Beyond descriptive and spatial evaluations, equity has also been incorporated directly into optimization models. Qian, Jaller, and Circella [41] developed a framework for the equitable distribution of bikeshare stations in Chicago, explicitly balancing efficiency objectives with spatial fairness. Their model demonstrated that relatively small adjustments in station siting could substantially reduce geographic inequities without imposing major efficiency losses. This perspective highlights

that equity is not only a policy concern but also an operational design variable that can be embedded into bikeshare planning and optimization.

Taken together, these studies show that equity in bikesharing requires attention to economic, spatial, and operational dimensions. For shared e-bike systems, equity concerns intersect with technical challenges such as battery availability and fleet distribution, reinforcing the importance of designing optimization strategies that account for both efficiency and fairness in system operations.

2.3 Operational Innovations and Optimization Foundations

As bikesharing systems have scaled, operational efficiency has become one of the most critical determinants of service quality. Systems must continuously balance bicycles across space and time, ensuring that users find available bikes and docks where and when they need them. This operational challenge has generated a large body of research on optimization models, routing strategies, and real-time management approaches.

One of the foundational studies is provided by Raviv, Tzur, and Forma [42], who formalized the problem of static repositioning in station-based systems. Their model addresses the overnight redistribution of bicycles using capacitated trucks to restore stations to target service levels. By formulating the problem as a mixed-integer program, they demonstrated that it is possible to achieve near-optimal rebalancing plans for medium-sized systems, while also proposing heuristic methods to handle larger networks. This work established the operational problem of bikeshare rebalancing as a variant of the vehicle routing problem (VRP), laying the groundwork for subsequent research on both static and dynamic repositioning.

Building on this foundation, Médard de Chardon, Caruso, and Thomas [43] examined real-world rebalancing practices across nine global cities. They found that rebalancing was often concentrated near transit hubs and in areas with strong directional flows, reflecting both commuter demand and operator priorities. Importantly, their analysis showed that operational strategies were not only about maximizing system efficiency but also about meeting policy objectives such as equity of access and service coverage. This highlighted the tension between profitability and social objectives

that continues to shape micromobility operations today.

More recent research has extended these models to account for demand dynamics and user incentives. Zhang and Meng [44] developed allocation and pricing strategies to encourage user participation in rebalancing, effectively crowdsourcing part of the operational burden. Their work illustrates a shift from purely operator-driven optimization to hybrid approaches that integrate user behavior, pricing mechanisms, and real-time information systems. Such strategies have become particularly relevant with the rise of dockless and electrified fleets, where flexibility and responsiveness are essential.

Together, these contributions illustrate the evolution of operational research in bikesharing: from static optimization models, to empirical analyses of rebalancing practices, to innovative strategies that incorporate demand management and user incentives. For shared e-bike systems, these insights remain crucial but require further extension. Battery management introduces additional constraints related to state-of-charge levels, vehicle capacity, and the temporal urgency of servicing depleted bikes. This thesis builds directly on the VRP and rebalancing literature by adapting and extending these models to the context of battery swapping, developing optimization frameworks that account for both spatial efficiency and energy availability in large-scale e-bike fleets.

Despite their popularity, shared e-bike systems present substantial operational challenges. One of the most critical is maintaining adequate battery availability across the fleet. Low state-of-charge (SOC) levels lead to decreased vehicle availability, reduced user satisfaction, and, ultimately, lower ridership levels [3]. For operators, this translates into revenue loss, higher complaint rates, and compromised reliability.

Operationally, battery management involves real-time monitoring, retrieval, redistribution, and in some cases direct charging of vehicles. These tasks are labor-intensive, time-sensitive, and expensive. As fleets expand, the scale of this challenge increases: even modest inefficiencies in battery logistics can propagate across the system, magnifying service disruption and operational cost.

Multiple strategies have been developed to manage battery availability in shared e-bike fleets:

Some systems equip docking stations with charging infrastructure, allowing e-bikes to recharge

when parked. This strategy takes advantage of natural usage cycles but requires substantial capital investment, including electrical upgrades and space for power equipment. In dense urban areas, such infrastructure may be infeasible due to space and permitting constraints.

Another approach is centralized depot charging, where depleted e-bikes are collected, transported to a depot, recharged, and then redeployed. While this avoids distributed infrastructure costs, it generates high operational expenses due to frequent vehicle movements and delays in reintroduction of bikes into the system.

Battery swapping offers a third pathway: depleted batteries are replaced with fully charged ones on-site, enabling near-instant redeployment. This method decouples charging from usage, reducing downtime and improving availability. Battery swapping is particularly advantageous in dockless systems, which lack fixed infrastructure and require mobile, flexible servicing solutions. However, swapping introduces a new set of logistical problems: managing battery inventory, determining which vehicles to prioritize, and planning van routes under capacity constraints.

The logistical challenge of battery swapping can be formally modeled as a vehicle routing problem (VRP), a classical optimization framework concerned with finding the most efficient routes for vehicles to deliver goods or services [45]. In the context of shared e-bike systems, the "goods" are charged batteries, and the service vehicle must select which bikes to visit, in what order, and under what resource constraints.

VRPs have been extensively studied across domains such as freight logistics, parcel delivery, and energy distribution. Variants relevant to battery swapping include capacitated VRPs, multi-depot VRPs, and VRPs with time windows. Each variant reflects specific operational realities—such as limited van capacity, multiple charging depots, or demand fluctuations—that parallel the needs of micromobility operators.

2.4 Vehicle Routing Problem (VRP) and Variants

The Vehicle Routing Problem (VRP) is a foundational problem in transportation and logistics optimization. It extends the classic Traveling Salesman Problem (TSP) by considering multiple vehi-

cles operating from one or more depots to serve a set of spatially distributed customers [45, 46, 47]. The goal is typically to minimize the total travel cost or distance while satisfying operational constraints such as vehicle capacity, time windows, or service requirements. Since its inception, numerous variants of the VRP have been developed to capture the diverse characteristics of real-world distribution and mobility systems. The most relevant variants to this research are described below.

The Capacitated VRP (CVRP) is the most classical and widely studied form of the VRP. In CVRP, each customer has a certain demand quantity, and each vehicle has a limited carrying capacity. The objective is to design the least-cost set of routes such that every customer is visited exactly once and the total demand on any route does not exceed the vehicle capacity [45]. The CVRP models resource limitations that are central to most logistics operations—for example, delivery trucks with limited cargo volume or, in micromobility contexts, service vans carrying a finite number of charged batteries. In the battery swapping scenario, each visit consumes one unit of capacity, representing the replacement of a depleted battery. The capacity constraint ensures that the number of swaps per trip does not exceed the number of charged batteries available at the start of the route.

The VRP with Time Windows extends the CVRP by introducing a temporal dimension. Each customer must be visited within a specified time interval, and early arrivals may require waiting until the service window opens [48, 49]. This variant captures operational constraints such as delivery schedules, charging station hours, or workforce shifts. In the context of e-bike battery swapping, time windows could represent operator working hours, restrictions on servicing residential areas, or demand peaks when bikes must be available at specific times of day. Solving VRPTW instances is significantly more complex than CVRP because of the added temporal feasibility dimension, often requiring specialized heuristics or branch-and-price algorithms.

The Pickup-and-Delivery VRP generalizes routing to scenarios where items must be moved between pairs of locations, each request consisting of one pickup node and one delivery node [50, 51]. This introduces precedence and load-balance constraints: each pickup must occur before its delivery, and vehicle load must be updated dynamically along the route. PDVRP is particularly relevant in applications like ride-sharing, parcel collection, and cargo transport. Battery swapping differs

slightly because the van both drops off and collects batteries at each stop, but there are no paired pickup-delivery relationships between different bikes. Nonetheless, PDVRP concepts remain useful when modeling more advanced operations such as multi-depot battery exchange or inter-hub redistribution.

The Prize-Collecting VRP (PCVRP) and the related Team Orienteering Problem (TOP) introduce the concept of optional visits, where each customer provides a profit or "prize" that is earned if the location is served, while travel incurs a cost [52, 53]. The objective is to maximize total profit minus travel cost, subject to limits on distance or vehicle availability. These models are ideal when not all customers need to be visited, an operator must select the most profitable subset to serve within resource constraints. In the e-bike swapping context, each bike yields a potential revenue gain from restoring its battery to a high state of charge, but distant or low-demand bikes may not justify the travel cost. The proposed methodology in this thesis aligns closely with this profit-based routing logic, combining selective service with operational efficiency.

The Selective VRP (also known as the Profitable VRP) generalizes the prize-collecting idea by explicitly modeling a fixed operational budget such as vehicle capacity, travel distance, or total route duration [54, 55]. Instead of serving all customers, the model selects a subset of nodes that maximizes the net gain. The SVRP provides a natural structure for the battery swapping problem because the service van can only perform a limited number of swaps within one cycle, and each potential visit has a different economic benefit depending on the bike's SOC and location. The decision variable x_i in the proposed formulation directly corresponds to the node-selection variable in the SVRP, marking whether bike i is worth visiting during the route.

In practical operations, multiple depots or vehicles are often used to cover large areas. The Multi-Depot VRP (MDVRP) assigns customers to depots and optimizes routes for several vehicles simultaneously [56, 57]. These extensions introduce an additional allocation layer and are relevant for city-scale battery swapping systems with several charging hubs or maintenance vans. While the present work focuses on the single-depot, single-vehicle case to establish methodological clarity, the proposed clustering and routing framework can be extended to the multi-depot case by treating

each cluster as an independent service zone anchored at its nearest depot.

At city scale, solving the VRP directly for thousands of nodes is computationally intractable. To address this, clustering methods are commonly used as a pre-processing step. The cluster-first, route-second (CFRS) paradigm partitions the problem into manageable subproblems, where vehicles service geographically or operationally coherent groups of nodes. This reduces computational complexity while preserving solution quality.

In battery swapping, clustering helps ensure that each van serves a compact region of the city, reducing travel distance and making capacity constraints easier to enforce. However, clustering must also account for SOC distribution: a spatially compact cluster with too many depleted bikes may exceed capacity, while a balanced cluster improves both feasibility and profitability.

San Francisco provides a particularly relevant testbed for this study. The Bay Wheels bike-share system, operated by Lyft, offers a large dockless e-bike fleet with real-time data available via the General Bikeshare Feed Specification (GBFS). The city's topography, high cycling demand, and heterogeneous trip patterns create diverse operational conditions, from flat downtown corridors to steep residential neighborhoods.

Furthermore, San Francisco's transportation policies emphasize sustainability and micromobility adoption, aligning closely with the goals of efficient battery management. By applying the proposed optimization framework to this setting, the study not only tests methodological robustness but also generates insights applicable to other dense North American cities facing similar challenges.

This chapter introduced the broader context of micromobility systems, highlighted operational challenges in e-bike fleets, and reviewed existing strategies for battery management. It also established the relevance of vehicle routing and clustering methods as foundational tools for solving large-scale optimization problems in this domain. Finally, it positioned San Francisco as a relevant case study, providing both practical motivation and empirical grounding for the methodology developed in subsequent chapters. Shared e-bike systems require operational strategies that balance user convenience with cost-effective fleet management. Battery availability, in particular, has emerged as a key challenge, shaping both service quality and economic viability. The growing body of literature

on charging, rebalancing, and routing optimization provides valuable insights, yet most studies have been tailored to either dock-based systems or simplified operational settings. This chapter reviews three relevant streams of literature: (i) battery swapping strategies in shared mobility, (ii) variants of the vehicle routing problem (VRP) that inform battery distribution logistics, and (iii) cluster-first, route-second (CFRS) methods designed to handle large-scale urban networks. Together, these perspectives establish the foundation for the optimization framework proposed in this thesis.

Battery swapping has emerged as a promising solution for maintaining high service availability in shared electric mobility. Unlike plug-in charging, which requires long idle times and fixed infrastructure, swapping allows operators to quickly replace depleted batteries on-site, reducing downtime and improving fleet utilization. Studies have examined various aspects of battery swapping, including station design, operational modes, and optimization strategies. For example, Wu [5] categorize battery swapping systems by decision scenarios, infrastructure types, and service levels, while Vallera et al. [58] review technological developments and highlight trade-offs between economic viability and operational complexity.

Operationally focused studies provide more granular models of swapping logistics. Xie et al. [8] examine locker-based deployments in urban e-bike networks, using queueing theory to reduce service delays and improve user access. Optimization-based approaches have also been explored. Shao et al. [6] propose a three-stage model combining clustering, routing, and vehicle assignment to maximize revenue and minimize operational distance. Reinforcement learning methods have been applied by Xu et al. [7], who employ a dueling DQN to make dynamic swapping and rebalancing decisions in New York City's dense network. Stochastic approaches have also been introduced, such as Markov Chain models to capture probabilistic energy depletion and replacement cycles [9].

Beyond academic research, real-world pilots demonstrate the practical feasibility of battery swapping. In Taiwan, the Gogoro scooter-sharing system has scaled battery swapping to hundreds of thousands of users, supported by dense networks of swap stations and logistics integration. In

China, large-scale pilots in both scooter and e-bike fleets have shown the potential for swapping to support high-frequency urban demand. These examples highlight the operational promise of swapping, though academic studies on e-bike systems remain primarily focused on modeling rather than implementation. Importantly, most existing work targets structured, dock-based systems, whereas dockless fleets pose unique challenges due to their diffuse spatial distribution and greater heterogeneity in state-of-charge (SOC).

The logistical problem of distributing charged batteries falls under the broader class of vehicle routing problems (VRPs), which are central to the field of operations research. Traditional VRP formulations involve designing minimum-cost routes for a fleet of vehicles to deliver goods from one or more depots to a set of customers. Numerous extensions have been developed to account for real-world constraints, including time windows, multi-trip operations, heterogeneous fleets, and partial deliveries. For example, Che et al. [59] address a multi-depot VRP with open inter-depot routes using a tabu-based adaptive large neighborhood search (T-ALNS), while Pan et al. [60] study a multi-trip time-dependent VRP with maximum route duration constraints, solved via a hybrid ALNS and VND algorithm.

In the context of electric mobility, energy-specific VRP extensions have attracted growing attention. Wang and Zhao [61] propose a partial charging strategy for electric vehicle routing, integrating energy consumption rates, charging times, and vehicle heterogeneity into the routing formulation. Similarly, Wang et al. [62] develop a rule-based framework for large CVRP instances that balances computational efficiency with route quality. These contributions illustrate the adaptability of VRPs to energy-constrained systems.

For micromobility, VRPs are particularly relevant to bike redistribution and rebalancing, where operators must move bicycles across a city to match spatially and temporally varying demand. Early works focused on static overnight rebalancing, while more recent studies incorporate dynamic demand and real-time decision making. Battery swapping extends this literature by combining features of inventory routing and rebalancing, since charged batteries must be distributed while depleted ones are simultaneously collected. This dual-flow nature makes battery swapping more

complex than conventional bike rebalancing, linking it closely to both logistics and energy delivery problems.

As city-scale e-bike networks expand, computational scalability becomes a critical concern. Solving a VRP directly for thousands of service nodes is intractable with exact methods, motivating decomposition techniques such as the cluster-first, route-second (CFRS) approach. CFRS reduces complexity by dividing the problem into two stages: bikes are grouped into clusters based on spatial or operational criteria, and routing is then solved locally within each cluster. This strategy has been widely applied in logistics and parcel delivery and has recently been adapted to micromobility.

For example, Shao et al. [6] apply a capacitated k-medoids clustering method to group low-SOC bikes into coherent clusters before solving a two-layer ALNS for routing. Xiao et al. [63] design a route-grouping heuristic that integrates clustering and optimization for large CVRP instances. More recently, Cai et al. [64] use hierarchical density-based clustering (HDBSCAN) to partition bicycle networks into service zones prior to routing. Reviews by Qi et al. [65] and Santos et al. [66] emphasize that decomposition frameworks and adaptive metaheuristics like LNS and ALNS are essential tools for handling real-world routing problems with thousands of nodes. Reinforcement learning also provides new scalability pathways, as shown by Xu et al. [7] in their work on dynamic bike rebalancing.

CFRS is particularly valuable for dockless e-bike systems, where spatial distribution is diffuse and state-of-charge varies widely across the fleet. Clustering reduces computational burden and ensures that routing decisions remain feasible within operational limits such as van capacity and service times. By combining clustering with local optimization, CFRS approaches bridge the gap between large-scale system complexity and practical routing efficiency.

Despite these advancements, several research gaps persist. First, much of the literature on battery management assumes station-based infrastructure or relies on simulated data, limiting its direct applicability to dockless systems. Second, optimization models often prioritize distance minimization while neglecting profitability, technician capacity, or user-level equity. Third, clustering strategies are frequently static and disconnected from real-time SOC information, reducing their

responsiveness in dynamic environments. Finally, few studies evaluate strategies across multiple urban contexts, leaving open questions about generalizability.

2.5 Research Gaps and Thesis Contribution

The literature on shared micromobility highlights both its rapid growth and the operational challenges that threaten its long-term sustainability. While e-bikes expand access and resilience compared to conventional cycling, their reliance on batteries introduces new logistical problems. Existing studies have examined charging infrastructure, depot-based operations, and battery swapping, but much of this work remains focused on structured, station-based systems. Dockless fleets, which are increasingly dominant in North America and China, pose unique challenges due to their diffuse spatial distribution, heterogeneous trip patterns, and dynamic battery depletion. Few studies to date have explicitly modeled these complexities in an optimization framework.

Research on vehicle routing problems (VRPs) offers valuable tools for addressing micromobility operations, but important limitations remain. Classical VRP formulations minimize travel distance or cost, whereas shared e-bike systems require balancing multiple objectives, including user availability, operator profitability, and equity of access. Moreover, most VRP extensions applied to micromobility address either static rebalancing or simplified demand conditions, often neglecting the dual-flow nature of battery swapping, where charged and depleted batteries must be managed simultaneously. This gap limits the direct applicability of existing routing solutions to the realities of e-bike fleets.

Clustering methods such as the cluster-first, route-second (CFRS) paradigm provide computationally feasible ways to scale routing to thousands of service nodes. However, existing clustering approaches are largely static and based on spatial proximity alone, without incorporating operational characteristics such as state-of-charge distribution or vehicle capacity. As a result, clusters may be geographically compact but operationally imbalanced, undermining efficiency and profitability in large-scale dockless systems.

This thesis addresses these gaps by developing an optimization framework for battery swapping in

dockless e-bike sharing systems that integrates dynamic clustering with VRP-based routing. Unlike prior studies that focus primarily on minimizing distance, the proposed framework pursues dual objectives of maximizing operator revenue and minimizing operational cost. It explicitly accounts for state-of-charge distributions, van capacity constraints, and the dual-flow nature of swapping. Furthermore, the framework is validated on real-world data from San Francisco's Bay Wheels system, ensuring empirical relevance. In doing so, the thesis contributes to both the transportation planning literature on shared micromobility and the operations research literature on large-scale, energy-constrained vehicle routing.

Having established the historical, policy, and operational context of shared micromobility, as well as the methodological foundations provided by the literature on battery swapping, vehicle routing, and clustering, the next chapter turns to the development of the proposed optimization framework. Chapter 3 details the mathematical formulation, decision variables, and constraints that define the battery swapping and routing problem in dockless e-bike fleets. It also describes the clustering approach used to reduce computational complexity and the dual objectives of minimizing operational costs and maximizing revenue. Together, these methodological elements build directly on the gaps identified in the literature review, providing a structured foundation for the empirical case study in San Francisco presented in Chapter 6.

Table 3: Summary of key impacts of bikesharing systems

Dimension	Benefits	Limitations / Challenges
Environment		
	• Reduces car trips and associated emissions [33]. • Supports climate and energy goals. • Complements public transit for sustainable travel.	• Redistribution and rebalancing generate additional emissions [27]. • Manufacturing and disposal of bikes and batteries affect life-cycle impacts.
Health		
	• Increases physical activity and active travel minutes [34]. • Extends life expectancy, prevents premature mortality [35]. • Net health gains outweigh risks from pollution and crashes [36, 37].	• Some substitution from walking reduces net gains. • Exposure to traffic crashes and air pollution.
Equity		
	• Expands access to cycling among new user groups [39]. • Targeted equity stations and subsidies improve accessibility [40].	• Persistent gaps in low-income and minority communities. • Unequal distribution of service coverage [41].
Safety		
	• Crash severity for bikeshare users is lower than for private cyclists ("safety in numbers") [30]. • System design (sturdy bikes, visible branding) improves safety outcomes.	• Risk remains higher than for walking or public transit. • Safety outcomes vary across cities and cycling cultures.

3 Methodology

This chapter presents the methodological framework developed to optimize battery swapping operations in dockless e-bike sharing systems. The formulation builds on the challenges identified in the background and literature review, translating them into a mathematical optimization problem that captures the dual objectives of minimizing operational costs and maximizing revenue. The chapter first defines the decision variables, parameters, and constraints that characterize the battery swapping and routing problem. It then introduces the clustering procedure used to partition the service network into computationally manageable regions while accounting for state-of-charge distributions and vehicle capacity. Finally, the solution approach is described, integrating clustering with a vehicle routing model to generate feasible and efficient service plans. Together, these methodological components establish a structured framework that will be applied to the San Francisco case study in Chapter 6.

3.1 Problem Statement

We formulate the battery swapping problem in dockless e-bike sharing systems as a variant of the capacitated vehicle routing problem (CVRP). The objective is to determine a profitable and operationally efficient route for a service van that replaces depleted e-bike batteries with fully charged ones. The problem integrates two objectives: (i) maximizing net profitability from swapping decisions and (ii) minimizing total travel distance. This dual-objective formulation reflects the trade-off between economic and operational efficiency in large-scale micromobility systems.

Battery swapping in dockless e-bike systems is naturally expressed as a vehicle routing problem with a service-selection decision at each node. The operator must choose a subset of candidate bikes to serve and sequence those visits subject to a limited inventory of charged batteries. If every candidate were mandatory, the problem reduces to a classical capacitated vehicle routing problem with one depot and customer demands. In practice, service is discretionary and profitability-driven, so the formulation must couple routing with a binary swap decision at each bike. Unlike the clas-

sical CVRP, where every customer must be visited exactly once, the proposed formulation allows the optimization model to decide which bikes are serviced during a given routing cycle, while other bikes may remain unvisited if they are not economically attractive.

The battery swapping problem formulated in this study is closely related to several well-known variants of the Vehicle Routing Problem (VRP), yet it also extends beyond them in important ways. Understanding these relationships helps position the model within the broader optimization literature and clarifies the methodological choices made in this thesis.

First, when every e-bike in the network is required to be served ($x_i = 1$ for all i) and the van's battery capacity limits the number of feasible stops, the problem simplifies to the *Capacitated Vehicle Routing Problem* (CVRP). In this classical variant, each customer must be visited exactly once, and the goal is to minimize total travel distance or cost while respecting vehicle capacity. In our context, each e-bike can be viewed as a customer with a unit demand representing one charged battery, and the van's load corresponds to the number of charged batteries it can carry from the depot. The CVRP therefore forms a special case of our model when service selection is not optional and profitability does not influence routing decisions.

Second, if the van has sufficient capacity to serve many or all candidates, but travel distance still carries a cost, the problem behaves similarly to the *Prize-Collecting Vehicle Routing Problem* (PCVRP) or the *Team Orienteering Problem* (TOP). In these variants, each node offers a certain reward or "prize" that is earned if the node is visited, while travel between nodes incurs a cost. The objective is to determine a subset of nodes that maximizes total profit (rewards minus travel costs). In the proposed model, the incremental revenue term $(R_i^F - R_i^C)$ plays the same role as a prize, representing the economic gain from replacing the battery of bike i. A route that visits bikes with higher expected gains but keeps total distance short will yield greater net profit, following the same logic as the PCVRP or TOP.

Third, the formulation can also be viewed as a specific instance of the *Selective Vehicle Routing Problem* (SVRP). In the SVRP, not all potential customers must be visited; instead, the operator selects an optimal subset of nodes that maximizes total profit or minimizes total cost under resource

constraints such as time, capacity, or distance budgets. Our binary decision variable x_i serves precisely this function, determining which bikes are worth servicing within the current routing cycle. The van's finite number of charged batteries, represented by constraint (7), acts as a resource budget that limits how many nodes can be served. The problem therefore combines route selection and node selection simultaneously, as in the SVRP, but with a domain-specific profit function tied to energy levels and usage forecasts.

It is also useful to distinguish the present formulation from the *Pickup-and-Delivery Vehicle Routing Problem* (PDVRP). In PDVRP, each request consists of two linked locations, a pickup node and a corresponding delivery node, connected by precedence and load constraints. By contrast, battery swapping involves independent service points: at each stop, the van simply exchanges one depleted battery for one charged battery. There is no dependency or precedence between different stops, and each visit consumes one unit of vehicle capacity rather than generating a pickup-delivery pair. This simplification allows the routing component to remain closer to a CVRP structure, while still accounting for profit-based node selection.

Overall, the proposed model can be interpreted as a *profit-weighted, capacity-constrained, selective routing problem* that unifies features of several VRP families. It extends the classical CVRP by introducing an economic objective that prioritizes profitable bikes instead of mandatory service. It shares with the prize-collecting and selective VRP families the idea of optional visits and a trade-off between distance and reward, yet it differs from them by integrating the state of charge (SOC) as a dynamic attribute influencing node profitability. The SOC-based decision process links the routing and revenue models, providing an operationally meaningful criterion for when and where to perform swaps. This combination of profit, energy state, and routing constraints defines a hybrid problem structure that, while related to established VRP variants, represents a novel extension tailored to the operational realities of dockless e-bike sharing systems.

We assume a single depot, a single service vehicle per cluster, and deterministic inputs for locations, distances, and state of charge. Each visit consumes exactly one charged battery and restores the bike to a high state of charge that supports additional rentals. Each bike can be visited at most once

within a planning cycle. Travel costs d_{ij} are computed from preprocessed geodesic or network distances that are time-invariant within the horizon. These assumptions match day-to-day operator practice for short planning horizons and allow solvable models that still capture the core trade-offs. The operator faces two objectives: to maximize the incremental revenue unlocked by swapping and to minimize the distance travelled. In this chapter, we adopt a weighted-sum scalarization, shown in (2).

Even the simplified special cases contained in our formulation are NP-hard. With $x_i \equiv 1$, the problem contains the travelling salesman problem as a subcase. With a fixed tour length budget and binary rewards it contains a knapsack-type selection. The full model that couples selection and routing inherits this hardness. This motivates the two-stage solution strategy introduced later, where clustering limits instance size and the routing subproblems are solved to good solutions within fixed time caps.

Two choices are central. First, the incremental revenue term $(R_i^F - R_i^C)$ ties service selection to measurable gains that align with operator goals. If detailed demand models are available, R_i^X can incorporate time-varying rental rates, peak multipliers, or penalties for unmet demand. Second, the capacity resource is modelled as battery inventory rather than vehicle time. If time windows or driver shifts become binding in a deployment, the formulation can be extended by adding service-time terms and classical time-window constraints. The methodology in Section 4 remains valid under these extensions.

The problem takes as input a set of bike locations and states of charge, a distance matrix $\{d_{ij}\}$, a per-minute rental rate r, and bike-level usage forecasts s_i used in (1). Locations and SOC can be obtained from GBFS feeds at snapshot times. Distances are precomputed via Haversine or a street network. Usage forecasts may come from simple historical averages or a learned model; the methodology is agnostic to the forecasting approach. This thesis employs two distinct study settings with complementary roles. The San Francisco Bay Wheels system is used as the primary case study to evaluate the proposed framework under real-world operating conditions using observed data. In contrast, a smaller Seattle-based example is used solely for illustrative and benchmarking

purposes, allowing controlled experimentation, parameter sensitivity analysis, and visualization of routing behavior. The Seattle example is not intended to represent an operational deployment, but rather to support methodological understanding and comparative analysis.

The remainder of the chapter specifies variables, parameters, objectives, and constraints; introduces a clustering mechanism that creates spatially compact, capacity-aware regions; and presents a solution strategy that couples clustering with mixed-integer routing inside a feedback loop. These components prepare the ground for the San Francisco case study in Section 6. The battery swapping problem formulated in this chapter integrates three decision components into a unified optimization framework: (i) which bikes to service, (ii) how to route the service vehicle, and (iii) how to evaluate the profitability of each potential swap. The model operates over a single planning cycle, during which locations, SOC values, and demand forecasts are treated as deterministic. The operator begins with a fixed inventory of fully charged batteries and must select a profitable subset of bikes to visit while ensuring that the resulting route is continuous, capacity-feasible, and operationally realistic.

To evaluate swapping decisions, the model uses a SOC-dependent revenue function (Equation 1), which links each bike's potential earnings to its battery level and expected usage time. This revenue signal is combined with travel costs derived from the distance matrix $\{d_{ij}\}$ to form the weighted-sum objective in Equation (2). The formulation therefore couples service selection and routing choices: a bike is visited only if the anticipated revenue gain outweighs the marginal cost of including it in the route.

The decision-making process is supported by five categories of inputs: (1) spatial data from the GBFS feed, (2) SOC measurements at the decision epoch, (3) a precomputed distance matrix, (4) usage forecasts that quantify expected rental duration, and (5) operational parameters such as battery thresholds and van capacity. Together, these inputs allow the model to represent the core trade-off faced by operators: whether the incremental value of restoring a bike to full charge justifies the cost of detouring the vehicle to reach it. The remainder of the chapter formalizes these elements through variable definitions, constraints, and the two-stage solution procedure.

3.1.1 System Description

The system under study represents a dockless e-bike sharing network operated over an urban service area. Each e-bike is equipped with an on-board battery of limited capacity and communicates its location and state-of-charge (SOC) to the operator through a real-time data feed. At any given decision epoch, the fleet consists of N active bikes, indexed by $i = 1, \ldots, N$, each associated with a geographic coordinate (x_i, y_i) and a measured SOC level $l_i \in [0, 100]$. These data are typically obtained from General Bikeshare Feed Specification (GBFS) snapshots, which provide second-level accuracy on bike status and battery level.

From an operational perspective, a battery swap involves an operator vehicle (the service van) visiting a bike in the field, removing its depleted battery, and replacing it with a fully charged one. The removed battery is later transported to a charging depot for recharging and reuse in subsequent cycles. The service van departs from a central depot carrying C charged batteries, where C represents the maximum carrying capacity determined by vehicle space and battery weight limitations. Each swap operation consumes one unit of this inventory and extends the bike's usable service time until its SOC again falls below the minimum operational threshold.

E-bikes are eligible for swapping based on their SOC level relative to two operational thresholds: the lower bound $L_{\min}$ and the upper bound $L_{\max}$. Bikes with SOC below $L_{\min}$ are considered depleted and prioritized for swapping, while those above $L_{\max}$ are considered sufficiently charged and excluded from service consideration. The intermediate range $[L_{\min}, L_{\max}]$ corresponds to bikes whose profitability for swapping is uncertain and depends on the trade-off between potential revenue and travel cost. The revenue associated with e-bike i if a swap is performed is denoted by R_i^F, while the projected revenue without swapping is R_i^C. A swap is executed only if $R_i^F > R_i^C$, ensuring that operational resources are directed toward profitable interventions.

We compute revenue using a SOC-proportional model that links expected usage to available battery charge:

$$R_i^X = r \cdot s_i \cdot \frac{l_i^X}{100}, \tag{1}$$

where $X \in \{F, C\}$ indicates whether the bike receives a fully charged battery (F) or retains its current state of charge (C). Here, r is the rental rate per minute, s_i is the forecasted ride duration, and l_i^X is the battery level under decision X.

Each bike i is therefore characterized by two revenue components: (i) the expected future revenue if its battery is replaced, denoted R_i^F, and (ii) the projected revenue if its battery remains unchanged, denoted R_i^C. These values depend on the local rental demand and the remaining SOC, which together determine how many rides the bike can support before running out of charge. The net gain from servicing a bike, $(R_i^F - R_i^C)$, represents the incremental value obtained by restoring it to full operational capacity. A swap is executed only if this net gain justifies the associated travel cost.

The service area is modeled as a fully connected graph $G = (V, E)$, where $V = \{0, 1, \ldots, N\}$ represents the set of nodes, with node 0 corresponding to the depot and nodes $1, \ldots, N$ corresponding to bikes. Each edge $(i, j) \in E$ is associated with a travel cost d_{ij}, representing the distance between nodes i and j. Distances can be calculated either using Euclidean coordinates or, more accurately, using geodesic distances (Haversine formula) derived from GPS coordinates. In the case study presented later, the geodesic metric is adopted to capture realistic urban travel distances without requiring a detailed street network model.

Operationally, each bike can be visited at most once during a routing cycle. The van departs from the depot, performs a sequence of battery swaps, and returns to the same depot after exhausting its load or reaching the end of the planned route. This single-route assumption corresponds to a single-vehicle daily operation or one shift of a battery swapping team. The framework can be generalized to multiple vehicles by extending the index set of routes and adding depot allocation constraints, but this study focuses on the single-depot case to isolate the effects of clustering and profitability thresholds.

Finally, the routing problem is solved over a static snapshot of the system, assuming bike positions and SOC levels remain constant during one decision period (typically one to two hours). This assumption aligns with real operational practices, where system states are updated frequently, and routing decisions are re-optimized periodically. The resulting static model captures the short-term

tactical planning stage of daily swapping operations, providing actionable schedules for field teams based on the most recent fleet conditions.

3.1.2 Decision Variables

The proposed optimization model uses three sets of decision variables that together define which bikes are serviced, how the service vehicle travels between them, and in what order those visits occur. Each variable has a direct operational interpretation within the battery swapping process.

- **Service decision** (x_i): a binary variable indicating whether e-bike i is selected for a battery swap during the current routing cycle. The value $x_i = 1$ means that the van will visit bike i and replace its battery; $x_i = 0$ means the bike is skipped. This variable captures the economic decision of whether the expected profit from swapping the battery exceeds the marginal travel cost to reach that bike. The vector $\mathbf{x} = (x_1, \ldots, x_N)$ therefore represents the chosen subset of bikes to service within the capacity limit.

- **Routing decision** (y_{ij}): a binary variable equal to 1 if the van travels directly from node i to node j, and 0 otherwise. This variable encodes the connectivity of the route across the network and forms the basis of the vehicle routing structure. Together, the variables $\{y_{ij}\}$ define the directed path that starts at the depot, visits a subset of bikes for swapping, and returns to the depot. Route continuity, capacity, and subtour elimination constraints are all expressed in terms of y_{ij}.

- **Visit sequence indicator** (t_i): a continuous variable that represents the position or order of node i within the route sequence. The depot is assigned $t_0 = 0$, and subsequent nodes are assigned increasing integer values along the tour. These sequence indicators are used in the Miller–Tucker–Zemlin (MTZ) formulation for subtour elimination [67], ensuring that the route remains a single continuous loop without disconnected cycles.

The joint use of $\mathbf{x}$, $\mathbf{y}$, and $\mathbf{t}$ allows the model to simultaneously determine both *which bikes to visit* and *in what order* to visit them. The binary service decisions $\mathbf{x}$ control the selection of nodes,

while the routing variables **y** and sequence indicators **t** establish feasible paths that satisfy vehicle capacity and continuity constraints. Together, these decisions provide a complete representation of the operational plan for one swapping route, linking economic priorities (through x_i) with logistical feasibility (through y_{ij} and t_i).

From an implementation perspective, the combination of node-selection and routing variables transforms the problem into a mixed-integer linear program (MILP). Each binary variable adds a combinatorial dimension, and the number of routing variables y_{ij} grows quadratically with the number of bikes N. This combinatorial growth highlights the computational difficulty of large-scale instances and motivates the decomposition and clustering strategies introduced later in Section 4.

3.1.3 Objective Function

The optimization problem seeks to maximize the overall profit generated from battery swapping operations during one service cycle. Profitability depends on two opposing factors: the economic benefit of restoring depleted bikes to full charge, and the operational cost incurred by traveling between them. In this context, revenue is modeled as an *opportunity cost*, representing the potential trip revenue that would be lost if a bike remains unavailable due to low battery levels. The objective function therefore balances economic benefit and operational efficiency as follows:

$$Z = \sum_{i=1}^{N} (R_i^F - R_i^C) x_i \; - \; \sum_{i=0}^{N} \sum_{j=0}^{N} d_{ij} y_{ij} \tag{2}$$

where Z is the total objective value. The first summation represents the total avoided loss in expected trip revenue achieved by performing battery swaps, while the second term penalizes the total distance travelled by the service van.

For each e-bike i, the difference $\left(R_i^F - R_i^C\right)$ measures the incremental opportunity cost avoided by swapping its battery compared to leaving it unchanged. Specifically, this term captures the expected trip revenue that would otherwise be lost if the bike remains unavailable due to insufficient charge. It accounts for the expected number of rides that the bike can support once recharged and the average rental rate per minute. Bikes with higher future demand or larger increases in

available energy yield larger positive values of $(R_i^F - R_i^C)$, making them more attractive for service. Multiplying by x_i ensures that this avoided opportunity cost is counted only for the bikes actually selected for swapping.

3.1.4 Constraints

The optimization model is subject to a set of logical, physical, and operational constraints that ensure feasibility of the solution. These constraints govern battery eligibility, route continuity, vehicle capacity, and the elimination of subtours. Together, they ensure that the resulting plan is both mathematically consistent and operationally realistic.

(1) Battery eligibility

$$l_i \geq L_{\min} - M x_i, \quad \forall i \tag{3}$$

$$l_i \leq L_{\max} + M(1 - x_i), \quad \forall i \tag{4}$$

These constraints link the binary decision variable x_i to the state of charge (SOC) of bike i, thereby governing its eligibility for battery swapping. The parameter $L_{\min}$ defines the minimum SOC threshold below which a bike is considered depleted, while $L_{\max}$ represents the upper SOC threshold above which a bike is regarded as sufficiently charged.

When $x_i = 1$, the first constraint relaxes the lower bound on l_i, allowing the SOC of bike i to fall below $L_{\min}$ and identifying it as a candidate for swapping. When $x_i = 0$, the same constraint enforces $l_i \geq L_{\min}$, preventing bikes with adequate charge from being selected. Similarly, the second constraint restricts the assignment of $x_i = 1$ to bikes whose SOC does not exceed $L_{\max}$, while bikes with higher SOC values are excluded from swapping decisions.

Together, these Big-M constraints encode a threshold-based decision rule in which the binary variable x_i reflects whether a bike's SOC violates predefined lower or upper bounds, without artificially altering the underlying battery levels. This formulation ensures that swapping decisions are driven by observed battery conditions rather than imposed battery states, allowing the optimization to

focus operational effort on bikes for which swapping yields meaningful system-level benefits.

(2) Route continuity

$$\sum_{j=1}^{N} y_{ij} = x_i, \quad \forall i \tag{5}$$

$$\sum_{i=1}^{N} y_{ji} = x_i, \quad \forall i \tag{6}$$

These constraints ensure that every bike selected for service ($x_i = 1$) has both an incoming and an outgoing connection in the route. Intuitively, if the van decides to visit a bike, it must arrive there from a previous node and depart to the next one. For nodes that are not selected ($x_i = 0$), no links are required. This pair of equations enforces logical route continuity between the selection and routing variables, preventing the optimizer from "serving" a bike that is not connected to the path.

(3) Capacity

$$\sum_{i=1}^{N} x_i \leq C \tag{7}$$

The service van can carry at most C fully charged batteries per trip, which directly limits the number of bikes that can be swapped. This constraint enforces that the number of selected bikes does not exceed the available battery inventory. The capacity C depends on vehicle design, battery size, and loading weight limits. In operational practice, this constraint corresponds to ensuring that each route can be executed within a single vehicle load without requiring mid-route recharging or depot returns.

(4) Routing logic

$$\sum_{j=1}^{N} y_{ij} \leq 1, \quad \forall i \tag{8}$$

$$\sum_{i=1}^{N} y_{oi} = 1 \tag{9}$$

$$\sum_{i=1}^{N} y_{io} = 1 \tag{10}$$

These constraints formalize the route's structure. Equation (8) ensures that the van departs from

each visited node at most once, preventing multiple departures from the same location. Equations (9) and (10) enforce that the van starts and ends its journey at the depot (node 0). Together, these rules guarantee that the resulting path forms a single continuous loop that begins and ends at the depot, visiting each selected bike exactly once.

(5) Subtour elimination (MTZ formulation)

$$t_j \geq t_i + 1 - M(1 - y_{ij}), \quad \forall i, j \ (i \neq j, i, j \neq 0) \tag{11}$$

The Miller–Tucker–Zemlin (MTZ) constraints [67] prevent the formation of subtours—disconnected cycles that visit subsets of bikes without returning to the depot. The variables t_i represent the visiting order of each bike, and M is a large constant (commonly $M = N$) that effectively disables the constraint when no direct link $y_{ij} = 1$ exists between nodes i and j. When $y_{ij} = 1$, the constraint enforces that bike j must appear later in the sequence than bike i, ensuring that the route remains a single connected path. Although other subtour elimination methods exist (e.g., flow-based or cut-set formulations), the MTZ approach is adopted here for its simplicity and compatibility with small to medium-scale MILPs.

(6) Decision variable domains

$$x_i, y_{ij} \in \{0, 1\}, \quad \forall i, j \tag{12}$$

$$t_i \geq 1, \quad \forall i \tag{13}$$

The binary domains of x_i and y_{ij} specify the combinatorial nature of the selection and routing decisions. The integer variable t_i is defined over positive values to represent the order of visits, with the depot typically assigned $t_0 = 0$. These domains ensure that the MILP formulation remains logically well-posed and interpretable.

Collectively, constraints (3)–(13) enforce the physical, logical, and operational realities of the battery swapping process. The model ensures that:

(a) only bikes with sufficiently low SOC are eligible for service,

(b) each serviced bike is properly connected within the route,

(c) the van's capacity of charged batteries is never exceeded,

(d) the route begins and ends at the depot, and

(e) the resulting path forms one continuous tour without subtours.

In larger systems, some of these constraints can be replaced or relaxed for computational efficiency. For example, subtour elimination can be handled by adding constraints dynamically within a branch-and-cut framework, and route continuity can be expressed using flow-conservation constraints instead of binary inequalities. Similarly, the capacity constraint can be extended to account for two-sided flows (delivery of charged and collection of depleted batteries) in multi-depot or round-trip settings. The formulation presented here balances realism and tractability for single-depot, single-vehicle operations.

4 Solution Methodology

The battery swapping optimization problem formulated in Section 3 is a mixed-integer program with binary selection and routing variables. Such problems are NP-hard, and their computational complexity increases exponentially with the number of nodes. For realistic micromobility fleets that may include hundreds of e-bikes, solving the full problem directly would be computationally infeasible within operational time limits. To address this challenge, the proposed approach adopts a two-stage hierarchical framework that combines spatial clustering and route optimization within an iterative feedback loop.

The first stage groups e-bikes into geographically compact and operationally balanced clusters, thereby decomposing the overall problem into smaller and more tractable subproblems. Each subproblem corresponds to a local service region that can be handled by a single vehicle or a single dispatch cycle. The second stage then solves a mixed-integer linear program (MILP) for each cluster to determine the optimal set of swap decisions and the corresponding vehicle route. After each iteration, the routing outcomes are fed back into the clustering process to refine spatial boundaries and SOC thresholds, allowing the system to gradually converge toward a globally efficient configuration.

This cluster-first route-second paradigm follows a well-established decomposition strategy in large-scale routing problems [68, 69, 70], but introduces a novel extension: a feedback-driven refinement process where operational indicators from the routing stage (e.g., profitability, SOC imbalance, and capacity utilization) inform subsequent clustering decisions. The approach balances model accuracy and computational efficiency while reflecting the dynamic, data-driven nature of micromobility operations.

4.1 Two-Stage Framework

The overall workflow is illustrated in Figure 4 and summarized algorithmically in Algorithm 1. The two stages, clustering refinement and routing optimization, interact iteratively until no further

improvement in total profit is observed.

Stage 1: Clustering refinement The first stage partitions the fleet of e-bikes into K clusters using an enhanced K-Means procedure that accounts for both geographic proximity and the distribution of state-of-charge (SOC). The goal is to form clusters that are spatially coherent and operationally feasible given the van's capacity C. After initial clustering, each cluster V_k is evaluated using a composite score S_k that combines spatial cohesion and SOC imbalance (Equations (15)–(16)). Clusters with high S_k values are refined by reassigning poorly fitting bikes to neighboring clusters, provided that doing so improves both local and global metrics and does not violate capacity constraints.

Unlike traditional K-Means, which minimizes only Euclidean distance to centroids, this method dynamically adjusts clusters based on operational criteria: how many low-SOC bikes each cluster contains, whether its potential swap demand exceeds available capacity, and how balanced the workload is across the fleet. This operationally informed clustering ensures that each MILP subproblem in the next stage remains feasible and meaningful. The refinement continues until either no further improvements occur or a predefined iteration limit $I_{\max}$ is reached.

Stage 2: Routing optimization Once clusters are defined, the second stage solves a mixed-integer linear program within each cluster to identify (i) which bikes to service and (ii) the sequence of visits that maximizes profit. The MILP is solved using the open-source CBC solver through the `PuLP` Python library, with a per-cluster time limit of 60 seconds. This limit ensures tractability and mimics real-time operational constraints, where decisions must be updated frequently based on new data. The optimization yields an updated set of SOC thresholds $(L_{\min}, L_{\max})$, reflecting the observed distribution of profitable swaps. These thresholds are then averaged across clusters to update global parameters for the next iteration of clustering.

The interaction between the two stages forms a feedback loop: clustering defines the spatial scope of routing, while routing outcomes inform the redefinition of clusters. This mutual adaptation allows the system to converge toward a stable configuration in which both spatial compactness and profitability are locally optimized. Convergence is declared when total system profit changes by less than a small tolerance ε across consecutive iterations or when a maximum number of global

iterations is reached. In practice, convergence is typically achieved within five to eight iterations, depending on fleet size and heterogeneity.

The algorithm begins by initializing thresholds $L_{\min}^{(0)}$ and $L_{\max}^{(0)}$ using empirical SOC distributions. A moderate number of clusters (K) is chosen based on fleet size and desired workload balance. For example, if the service van can handle 25 bikes per route and the fleet contains 300 bikes, then $K = 10$ clusters provide a reasonable partitioning.

The computational cost of each iteration is primarily driven by two components. First, the clustering refinement stage runs in $O(NK)$ time, meaning that its runtime grows proportionally with the product of the number of data points and the number of clusters. This reflects the standard K-Means update, where each point is compared to each cluster center during reassignment. Second, the routing stage requires solving a separate MILP for each cluster. Although MILP complexity is exponential in the size of the decision set, each problem remains tractable because the number of bikes assigned to any cluster, denoted $|V_k|$, is intentionally kept small (at most 35). Since each cluster's routing problem is independent, all MILPs can be executed in parallel. As a result, the iterative procedure converts one large, computationally intractable MILP into multiple smaller and parallelizable subproblems, significantly improving overall scalability.

One main parameter influences the algorithm's behaviour. The weight α in Equation (14) balances spatial compactness against SOC balance during clustering. This parameter can be tuned through sensitivity analysis or calibrated using empirical operational data. A default value of $\alpha = 0.5$ has been found to yield stable convergence in preliminary experiments.

Because both clustering and routing stages are bounded optimization procedures, the iterative loop is guaranteed to converge to a locally optimal configuration with respect to the profit metric. However, the non-convex nature of the overall problem means that global optimality cannot be guaranteed. Empirical results show that initializing with multiple random seeds for the clustering centroids helps avoid poor local minima and yields consistent system-level outcomes.

From an operational standpoint, this two-stage framework mirrors how micromobility operators plan daily activities: first dividing the city into manageable service zones, then optimizing each

Algorithm 1 Battery Swapping Optimization with Feedback-Driven Clustering

Require: $\{x_i, l_i\}_{i=1}^{N}$ (bike locations and SOC), number of clusters K, truck capacity C, buffer, radius R, weight α
Ensure: Final clusters $\{V_k\}_{k=1}^{K}$, routes, and total profit

Initialize $L_{\min}^{(0)}, L_{\max}^{(0)}$, set $t \leftarrow 0$
$\{o_k^{(0)}\}_{k=1}^{K} \leftarrow$ K-Means$(\{x_i\})$
Assign each bike B_i to nearest o_k subject to $|V_k| \le C +$ buffer
Partition each cluster into M_k, D_k, N_k using $L_{\min}^{(t)}, L_{\max}^{(t)}$
repeat ▷ outer loop
  Stage 1: Clustering refinement
  repeat
    for $k = 1$ to K **do**
      Compute c_k and SOC weight w_k (Eq. (14))
      Compute imbalance H_k and cluster score S_k (Eq. (14))
    end for
    $k^{\text{worst}} \leftarrow \arg\max_k S_k$
    Select worst bike $B_j^{\text{worst}} \in V_{k^{\text{worst}}}$ using bike score $S_b(\cdot)$ (Eq. (18))
    for all neighbor clusters m with $d(o_{k^{\text{worst}}}, o_m) \le R$ **do**
      if $|V_m| < C +$ buffer and reassignment improves scores **then**
        Move B_j^{worst} from $V_{k^{\text{worst}}}$ to V_m
        Update centers $o_{k^{\text{worst}}}$ and o_m
        break
      end if
    end for
  until no improving move found or max inner iterations reached
  Stage 2: Route and swap optimization
  for $k = 1$ to K **do**
    Solve cluster MILP (time cap 60 s)
    Obtain $L_{\min,k}^{\star}$ maximizing cluster profit
  end for
  $L_{\min}^{(t+1)} \leftarrow \frac{1}{K} \sum_{k=1}^{K} L_{\min,k}^{\star}$
  $L_{\max}^{(t+1)} \leftarrow L_{\max}^{(0)}$
  Recompute M_k, D_k, N_k for all clusters
  $t \leftarrow t + 1$
until $L_{\min}^{(t)}$ converges or $t = t_{\max}$
return $\{V_k\}_{k=1}^{K}$, routes, and total profit

crew's route based on up-to-date SOC information. The algorithm thus provides not only a mathematical optimization tool but also a decision-support process that aligns naturally with the spatial and temporal structure of real-world e-bike operations.

4.2 Clustering Refinement Algorithm

The first stage of the solution methodology aims to partition the fleet into spatially coherent and operationally balanced clusters. This step dramatically reduces the dimensionality of the routing

problem by decomposing the full service area into smaller subproblems that can each be solved by a single vehicle. However, a purely distance-based clustering such as standard K-Means tends to produce regions of uniform geographic size without considering operational factors such as SOC imbalance or capacity constraints. To overcome this limitation, an enhanced clustering refinement algorithm is introduced that explicitly accounts for both spatial compactness and energy distribution.

In practice, not all parts of the city require equal service effort. Some areas may contain dense clusters of low-SOC bikes requiring immediate swaps, while others may have mostly charged bikes. A uniform K-Means clustering would treat these zones equally, leading to unbalanced workloads across clusters. The proposed refinement step corrects these imbalances by dynamically reassigning bikes between clusters based on a composite performance score that blends spatial and energy-based criteria. This ensures that each cluster remains both geographically compact and operationally feasible given the service vehicle's battery capacity.

Each cluster V_k is evaluated through a composite score S_k defined as:

$$S_k = \alpha c_k + (1-\alpha) H_k, \tag{14}$$

where c_k measures the spatial dispersion of bikes around the cluster centroid, H_k quantifies the imbalance in swap demand across clusters, and $\alpha \in [0,1]$ controls the relative importance of spatial versus energy-based criteria.

Spatial cohesion is computed as the average Haversine distance between each bike and the cluster center:

$$c_k = \frac{1}{|V_k|} \sum_{B_j \in V_k} d(B_j, o_k), \tag{15}$$

where $d(\cdot)$ is the great-circle distance and o_k denotes the cluster centroid. Smaller c_k values indicate tighter, more compact clusters with shorter average travel distances.

SOC imbalance H_k captures how far the number of low-SOC bikes in cluster k deviates from the

fleet-wide average, defined as:

$$H_k = \left| w_k - \frac{1}{K} \sum_{k'=1}^{K} w_{k'} \right|, \tag{16}$$

where w_k is a weighted measure of swap demand in cluster k:

$$w_k = |M_k| + \sum_{B_j \in D_k} \frac{L_{\max} - l_j}{L_{\max} - L_{\min}}. \tag{17}$$

Here, M_k, D_k, and N_k represent subsets of bikes classified respectively as must-swap, decision-required, and no-swap. The weighting term smoothly scales partial swap demand in proportion to how close a bike's SOC is to the lower threshold $L_{\min}$. This formulation avoids abrupt binary classifications and better reflects gradual energy depletion across the fleet.

Within each cluster, individual bikes are also evaluated using a bike-level score $S_b(B_j)$ that measures how well a bike fits its current cluster in both spatial and operational terms:

$$S_b(B_j) = d(B_j, o_k) + \lambda u_j, \tag{18}$$

where $d(B_j, o_k)$ measures spatial deviation from the cluster centroid and u_j quantifies the contribution of bike j to SOC imbalance. The second component u_j is defined as:

$$u_j = \begin{cases} -\dfrac{w_k}{|V_k|}, & \text{if } B_j \in N_k, \\ \dfrac{L_{\max} - l_j}{L_{\max} - L_{\min}} - \dfrac{w_k}{|V_k|}, & \text{if } B_j \in D_k, \\ 1 - \dfrac{w_k}{|V_k|}, & \text{if } B_j \in M_k. \end{cases} \tag{19}$$

Intuitively, u_j captures whether bike j increases or decreases the local workload imbalance. A bike far from its centroid or belonging to an overloaded cluster (high w_k) will have a higher $S_b(B_j)$, making it a candidate for reassignment to a neighboring cluster.

The refinement proceeds iteratively:

(i) Compute S_k for all clusters.

(ii) Identify the cluster with the highest S_k (worst-performing cluster).

(iii) Within that cluster, identify the bike with the highest $S_b(B_j)$.

(iv) Evaluate reassignment of this bike to neighboring clusters based on the marginal change in both S_k and capacity feasibility.

(v) Accept the reassignment if it decreases both the bike-level and global scores without violating the capacity constraint $|V_k| \leq C$.

After each accepted move, cluster centroids o_k and swap weights w_k are recalculated. The process continues until no improvement is observed or the iteration limit $I_{\max}$ is reached.

Consider three clusters, each initially containing 30 bikes. Cluster 1, located near the city center, has 18 low-SOC bikes, while Cluster 2 has only 5. The global average number of low-SOC bikes is therefore 10 per cluster. Cluster 1's imbalance component is $H_1 = |18 - 10| = 8$, while Cluster 2's is $H_2 = |5 - 10| = 5$. If a low-SOC bike near the boundary between clusters 1 and 2 is reassigned from Cluster 1 to Cluster 2, both H_1 and H_2 move closer to the average, reducing total imbalance. Provided that the receiving cluster does not exceed capacity and the new bike's spatial distance to its new centroid is smaller or comparable, the reassignment is accepted. Over successive iterations, such local adjustments produce clusters that are not only compact but also balanced in terms of swap workload.

The weighting parameter α in Equation (14) plays a key role in balancing spatial and operational objectives. A higher α places more emphasis on geographic compactness, producing tightly bound clusters at the cost of possible workload imbalance. A lower α prioritizes energy balancing, potentially enlarging clusters but ensuring fairer distribution of depleted bikes. Similarly, the coefficient λ in Equation (18) controls the relative penalty for imbalance versus distance. Operators can tune these parameters empirically based on priorities such as minimizing travel distance, equalizing workload among crews, or stabilizing SOC distribution across the network.

Compared to standard K-Means, this refinement algorithm offers three key advantages:

(a) It enforces operational feasibility by incorporating vehicle capacity into clustering decisions.

(b) It integrates SOC information directly into cluster formation, making clusters meaningful for energy-based routing.

(c) It allows iterative, local adjustments that converge quickly and do not require solving additional optimization problems at each step.

This hybrid structure enables efficient decomposition of large-scale instances while maintaining alignment between clustering logic and operational goals. The refined clusters serve as the input for the routing optimization stage described in Section 4.3.

4.3 Route and Swap Optimization

Once the clusters are refined, the second stage of the methodology determines the optimal subset of bikes to service and the route the van should follow within each cluster. This stage directly applies the mixed-integer linear program (MILP) formulated in Section 3, using the decision variables, objective function, and constraints previously defined. Each cluster is treated as an independent routing subproblem, with its own set of bikes, distances, and local parameters.

The goal of the MILP within each cluster is to jointly optimize service selection (x_i) and routing decisions (y_{ij}) so as to maximize total profit while respecting the van's capacity and route feasibility constraints. The optimization weighs the incremental revenue gained from swapping a bike's battery, $(R_i^F - R_i^C)$, against the travel cost associated with reaching that bike. The result is a set of binary decisions:

- $x_i = 1$ if bike i is selected for swapping;
- $y_{ij} = 1$ if the route includes a direct move from i to j;
- t_i indicating the order of visits.

Together, these variables define a feasible service route that maximizes the cluster-level profit. Because the model includes both selection and routing, it simultaneously solves a variant of the

Selective Capacitated Vehicle Routing Problem, where only a subset of candidate nodes are visited for maximum net gain.

The routing optimization is implemented in Python using the open-source library `PuLP`, which provides a convenient modeling interface for linear and mixed-integer programs. The MILP is solved using the `CBC` (Coin-or Branch and Cut) solver. CBC is chosen for its open-source availability, stability, and good performance on medium-sized routing instances. Each cluster-level MILP is solved independently, with a solver time limit of 60 seconds to maintain tractability during iterative feedback. This limit mimics real-world operational settings, where decisions must be refreshed frequently as new SOC data become available. Preliminary tests showed that most clusters (with 20–35 bikes) reach near-optimal solutions well within this time frame, with an optimality gap below 3–5%.

Within each cluster, the MILP ensures that the route begins and ends at the depot associated with that region, typically the same as the global depot for the single-vehicle case. The capacity constraint ($\sum x_i \leq C$) ensures that the number of battery swaps on the route never exceeds the number of charged batteries carried by the van. The MTZ constraints maintain route connectivity, eliminating subtours and enforcing a single continuous loop. In practice, each MILP subproblem yields a feasible tour even if not perfectly optimal, ensuring that every iteration of the overall algorithm remains valid for operational implementation.

After the MILP is solved for all clusters, the results are analyzed to update the SOC thresholds $L_{\min}$ and $L_{\max}$. Specifically, the algorithm records the minimum and maximum SOC levels among the bikes that were actually selected for swapping. These observed thresholds reflect the empirical boundaries between profitable and unprofitable swaps. By averaging these values across all clusters, the framework establishes new global thresholds for the next iteration:

$$L_{\min}^{(t+1)} = \frac{1}{K}\sum_{k=1}^{K} L_{\min}^{(k)}, \quad L_{\max}^{(t+1)} = \frac{1}{K}\sum_{k=1}^{K} L_{\max}^{(k)}.$$

This feedback mechanism allows the system to adapt to evolving SOC distributions and profitability conditions, aligning the decision thresholds with observed outcomes rather than relying on static

predefined values.

The iterative loop between clustering and routing continues until convergence is achieved. Convergence is defined as one of the following conditions:

(a) The total profit across all clusters changes by less than a specified tolerance ($\varepsilon = 0.5\%$) between consecutive iterations;

(b) No bike reassignments occur during clustering refinement;

(c) A maximum number of iterations $I_{\max}$ (typically 10) is reached.

If any of these conditions hold, the algorithm terminates and the most recent configuration of clusters and routes is recorded as the final solution. The output includes:

- The optimized route for each cluster, represented as a sequence of bike indices;
- The set of bikes selected for swapping ($x_i = 1$);
- The corresponding travel distances, total cost, and profit for each route;
- The updated SOC thresholds $L^*_{\min}$ and $L^*_{\max}$.

The routing optimization stage provides not only optimal or near-optimal service routes but also valuable operational insights. The proportion of selected bikes ($\sum x_i / N$) indicates the effective service coverage rate, while the average distance per swap ($\sum d_{ij} y_{ij} / \sum x_i$) reflects the efficiency of field operations. The evolution of SOC thresholds across iterations reveals the natural profitability range for battery replacements under current fleet conditions. These outputs can be used by operators to plan staffing levels, estimate energy logistics requirements, and design fair workload distributions among service teams.

Because each cluster-level MILP is independent, the routing optimization stage is inherently parallelizable. In practice, separate clusters can be solved simultaneously on multiple CPU cores or distributed computing nodes. This property enables the method to scale efficiently to fleets with

thousands of bikes while maintaining practical runtime. The modular structure also allows integration with commercial solvers such as Gurobi or CPLEX if higher performance is needed, without altering the formulation itself.

In summary, the routing and swap optimization stage transforms the outputs of the clustering process into actionable operational plans. Each MILP subproblem produces a feasible, capacity-constrained route that maximizes local profit, while the feedback mechanism continuously refines global parameters such as SOC thresholds. Together, these components form a self-adaptive optimization framework that balances spatial efficiency, profitability, and computational feasibility for large-scale e-bike battery swapping systems.

Algorithm 1 and Figure 4 summarize the proposed two-stage optimization framework. The algorithm alternates between clustering refinement and routing optimization, updating parameters after each cycle until the system converges to a stable configuration. This section provides a detailed explanation of how each component operates and how data flows through the process.

Initialization (Steps 1–3) The algorithm begins by importing the operational data, typically obtained from the General Bikeshare Feed Specification (GBFS) feed, which includes bike locations, states of charge (SOCs), and timestamps. These inputs are preprocessed to filter out inactive or missing bikes and to compute travel distances between all pairs of nodes, forming the distance matrix d_{ij}. The parameters for the first iteration, namely, the number of clusters K, vehicle capacity C, and SOC thresholds $L_{\min}^{(0)}$ and $L_{\max}^{(0)}$, are initialized based on historical data or operator experience. This initialization defines the feasible operational scope before optimization begins.

Stage 1: Clustering refinement In the first stage, e-bikes are partitioned into K clusters using K-Means applied to geographic coordinates. Each cluster represents a preliminary service zone that will later be refined based on operational conditions. Within each cluster, bikes are categorized into three classes:

- *Must-swap (M_k)*: bikes with $l_i < L_{\min}$, requiring immediate service;
- *Decision-required (D_k)*: bikes within the intermediate SOC range, where swapping profitability is uncertain;

- *No-swap (N_k)*: bikes with $l_i > L_{\max}$, which are sufficiently charged.

After classification, each cluster is evaluated using the composite score S_k (Equation (14)), which combines spatial compactness (c_k) and SOC imbalance (H_k). Clusters with high scores are considered inefficient either because they are spatially dispersed or because they contain too many (or too few) depleted bikes relative to other clusters. The algorithm identifies the worst-performing cluster and attempts to improve it by reassigning its least compatible bikes, as measured by $S_b(B_j)$ (Equation (18)), to neighboring clusters. A reassignment is accepted only if it decreases both local and global scores while respecting the van's capacity constraint. The refinement process repeats iteratively until no improving reassignments are found or a predefined iteration limit is reached.

Stage 2: Routing optimization Once the clusters are refined, the algorithm proceeds to the second stage. Each cluster is treated as an independent routing subproblem, formulated as a mixed-integer linear program (MILP) using the model defined in Section 3. The MILP determines which bikes to service (x_i), in what sequence (t_i), and through which direct connections (y_{ij}) the van travels. The objective function (Equation (2)) maximizes total net profit by balancing additional revenue from swaps against the travel cost incurred. The optimization for each cluster is solved with the CBC solver, subject to a 60-second time cap per instance. These cluster-level optimizations are independent and can be executed in parallel, which significantly improves scalability.

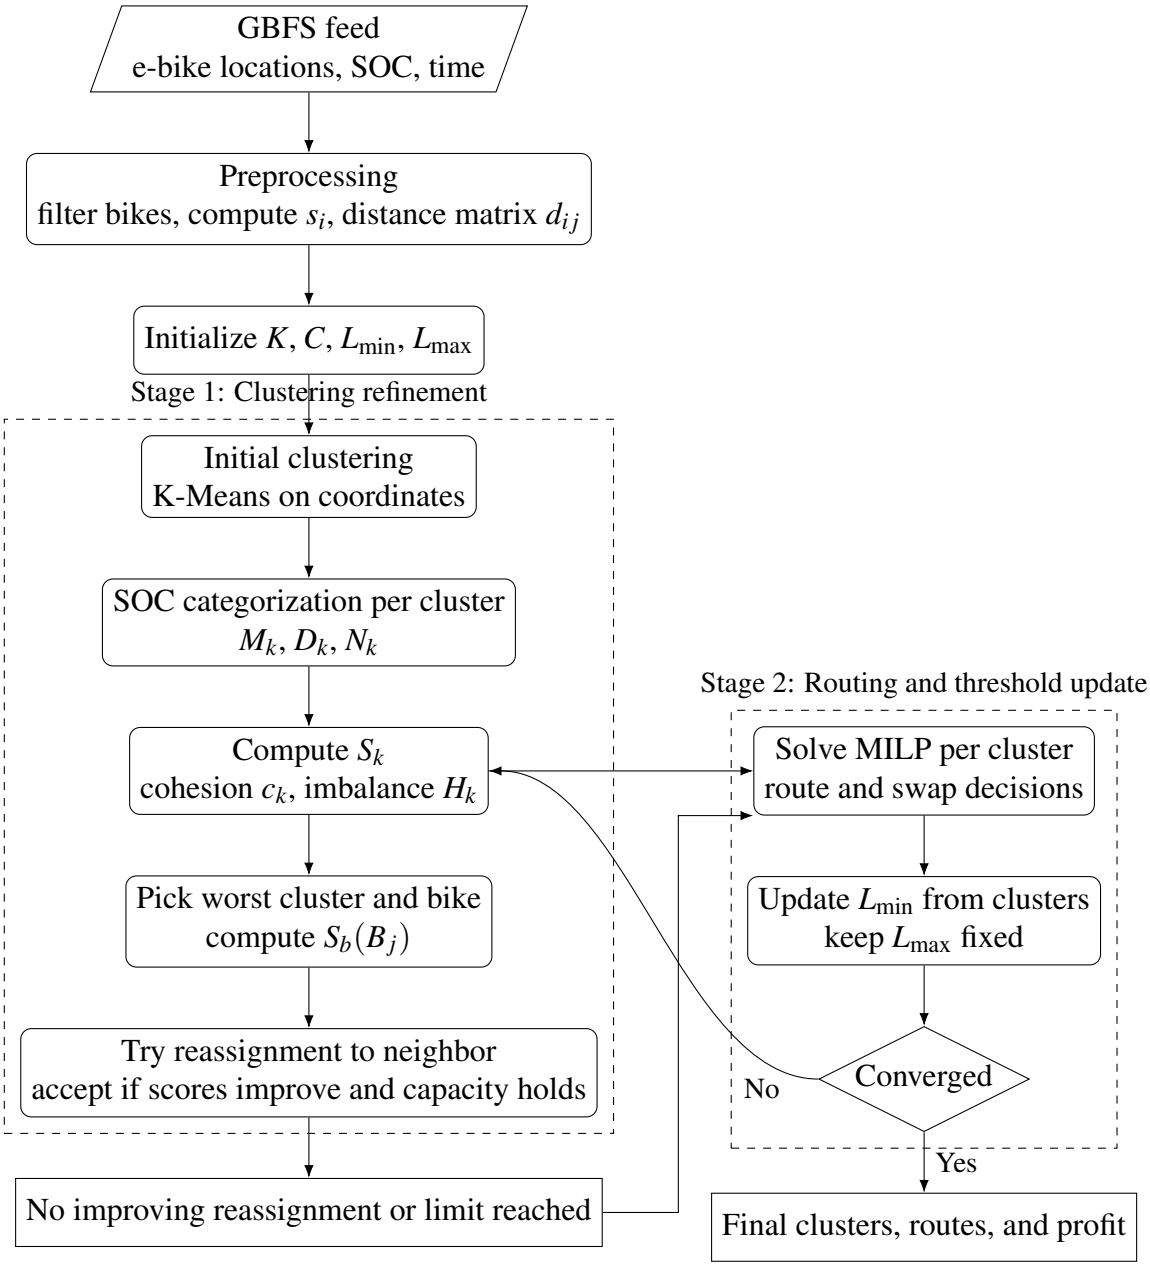

Figure 4: Two-stage workflow. Stage 1 forms and refines clusters using spatial and SOC information. Stage 2 solves a MILP per cluster and feeds updated thresholds back to Stage 1 until convergence.

5 Results

To evaluate the robustness and generalizability of the proposed battery swapping framework, a series of sensitivity analyses was conducted. These tests assess how variations in key parameters—such as the state-of-charge (SOC) thresholds and clustering weight—affect overall system performance. Sensitivity analysis is crucial in operational research models because input parameters often reflect uncertain or controllable operational policies rather than fixed constants. In the context of shared e-bike systems, changes in SOC thresholds or clustering priorities can significantly alter route efficiency, workload balance, and profitability. The following subsections present two complementary analyses: (1) the effect of SOC threshold policies on routing and profitability, and (2) the influence of the clustering weight parameter on spatial compactness and operational balance.

5.1 Sensitivity Analysis on SOC Threshold Policies

Battery state-of-charge (SOC) thresholds play a central role in determining which bikes are selected for battery swapping. In practice, these thresholds translate to operational policies that influence not only system profitability but also labor effort, vehicle mileage, and long-term battery health. To evaluate the effect of different SOC policies, we conduct a sensitivity analysis comparing single-threshold benchmarks with the dual-threshold model introduced in Section 3. The analysis highlights how threshold design interacts with van capacity and swap eligibility rules to shape system performance.

Table 4 summarizes the performance of four benchmark models (B1–B4) and three dual-threshold models (Original1–Original3) under varying van capacities. All models optimize the profit-minus-distance objective defined in Equation 2, balancing revenue from swaps with the travel cost of servicing.

The benchmarks illustrate the limitations of single-threshold policies. For example, B1 enforces swaps for all bikes with SOC below 30%, saturating the van capacity with 15 swaps but yielding a

Table 4: Performance metrics for benchmark and original models

Model	$L_{\min}$	$L_{\max}$	C	Objective	Profit	Distance (km)	Swaps
B1: force $L_{\min}$	30%	–	15	1357.1	1362.0	4.87	15
B2: force $L_{\min}$	20%	–	15	1069.2	1077.8	4.28	7
B3: force $L_{\min}$, allow others	20%	–	15	1395.5	1400.6	5.05	15
B4: force $L_{\min}$, allow others	20%	–	45	2362.7	2369.4	6.72	45
Original1	20%	80%	15	1395.5	1400.6	5.05	15
Original2	20%	80%	30	1913.0	1919.0	6.00	30
Original3	20%	80%	45	2223.4	2230.4	6.96	40

relatively low objective. Reducing $L_{\min}$ to 20% in B2 makes fewer bikes eligible (7), reducing both distance and profit. Allowing swaps beyond the minimum, as in B3, increases profit, but at the cost of including bikes that may not need immediate service. Scaling capacity further in B4 maximizes profit but at the expense of visiting all 45 bikes, inflating distance and workload.

By contrast, the dual-threshold models introduce an upper bound $L_{\max} = 80\%$, preventing unnecessary swaps on healthy batteries. Original1 achieves the same objective as B3 with $C = 15$, but avoids redundant interventions. As capacity rises, the gap widens: Original3 performs 40 swaps compared to 45 in B4, achieving nearly the same profit (2230.4 vs. 2369.4) but with a shorter route. This demonstrates that enforcing $L_{\max}$ improves efficiency and resource targeting without major loss of profit. These benchmark comparisons reveal not only the quantitative performance differences but also the underlying behavioral mechanisms of the system. Single-threshold models behave conservatively, prioritizing low-SOC bikes without considering redundancy or diminishing returns. This often leads to over-servicing or inefficient use of van capacity. In contrast, the dual-threshold models introduce a more selective logic: by capping $L_{\max}$, they effectively filter out bikes that would contribute little to system revenue relative to their travel cost. This mirrors a real-world operational strategy where swapping teams focus on high-impact interventions rather than uniformly servicing all low-battery bikes.

To better illustrate how van capacity influences routing structure beyond the numerical trends reported in Table 4, Figure 5 compares the optimized routes obtained for capacities of $C = 15$ and $C = 30$ under identical SOC conditions. In this illustrative example, initial battery state-of-charge

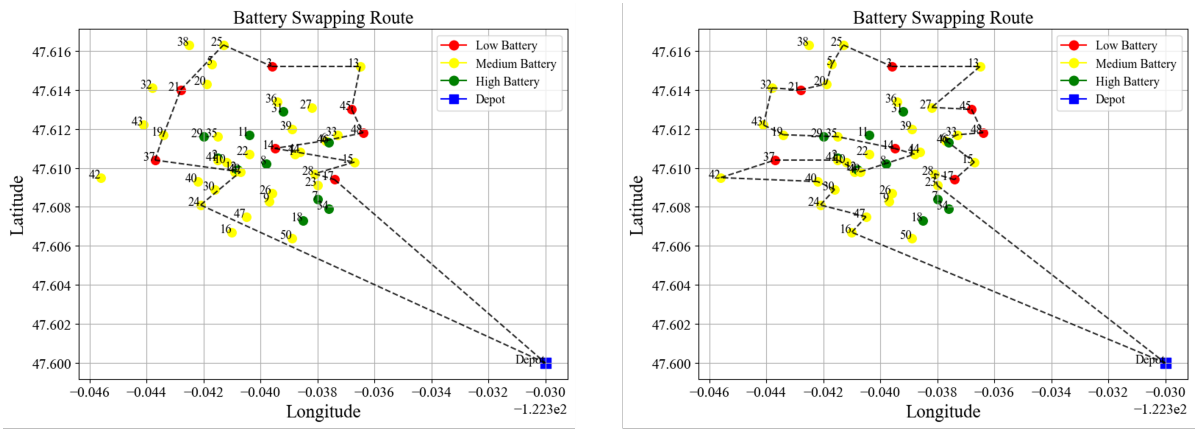

a) **b)**

Figure 5: Comparison of optimized battery-swapping routes for two van capacities. The left panel shows the route for capacity $C = 15$, and the right panel shows the route for $C = 30$. Red, yellow, and green markers denote low-, medium-, and high-SOC bikes, and the blue square represents the depot. Initial SOC levels in this example are randomly generated to illustrate heterogeneous battery conditions, and bike locations are based on an illustrative spatial layout from the city of Seattle. Increasing capacity reduces the need for intermediate returns and enables longer continuous tours.

values are randomly generated, and bike locations are drawn from a Seattle-based spatial configuration, in order to demonstrate routing behavior under heterogeneous conditions. With a smaller vehicle, the route contains shorter loops and more frequent returns toward the depot region, since the van quickly reaches its swap limit. In contrast, the route for $C = 30$ forms longer continuous tours and covers a larger spatial extent before returning, reflecting the additional capacity available for collecting medium- and low-SOC bikes. Although higher capacity increases total distance traveled, it also enables the system to capture more profitable swaps within a single outing. The visual comparison highlights the operational trade-off: larger vehicles improve revenue potential but produce longer and more complex tours, while smaller vehicles enforce a naturally compact service pattern.

Figures 6 and 7 illustrate how varying $L_{\min}$ affects performance under two policies: (i) "Allow," where bikes above $L_{\min}$ may still be swapped if capacity allows, and (ii) "Force," where only bikes below $L_{\min}$ are eligible.

Under the Allow policy (Figure 6), profit and distance remain stable up to $L_{\min} \approx 40\%$, after which they drop sharply as too few bikes remain eligible. Larger van capacity ($C = 30$) consistently yields

higher profit but also longer tours.

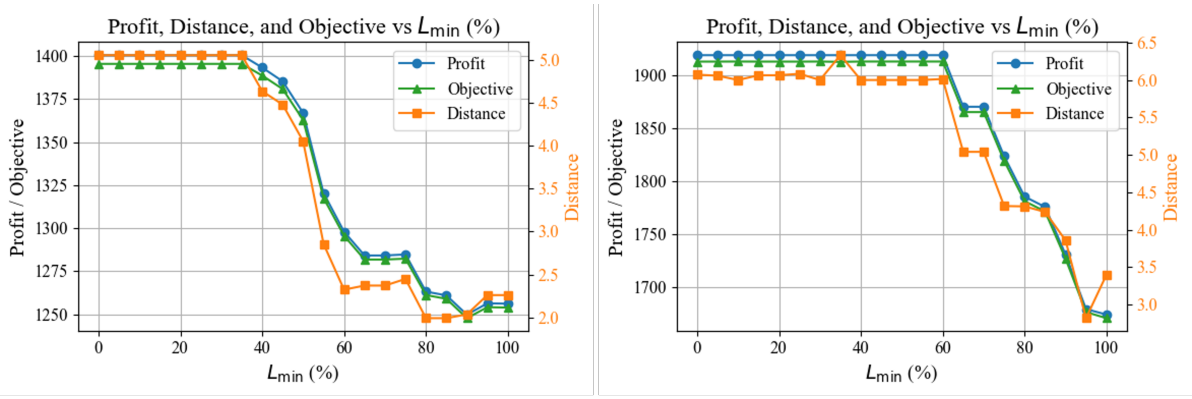

Figure 6: Sensitivity to $L_{\min}$ under the Allow policy for $C = 15$ and $C = 30$.

By contrast, the Force policy (Figure 7) shows rising profit as $L_{\min}$ increases up to about 40–50%, since the van focuses on low-SOC bikes that generate higher marginal returns. Beyond this point, eligible bikes become scarce, reducing profit and swaps. The larger capacity again provides higher absolute profit but also greater exposure to diminishing returns at strict thresholds.

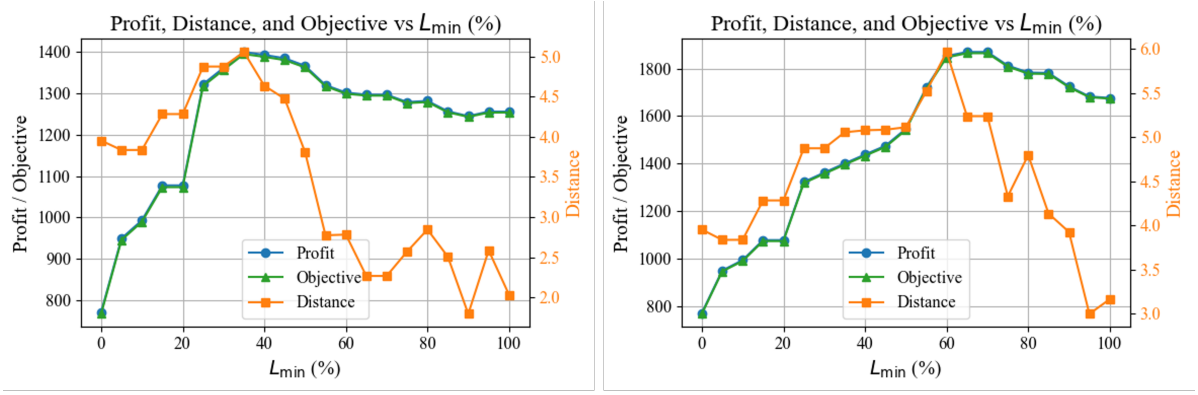

Figure 7: Sensitivity to $L_{\min}$ under the Force policy for $C = 15$ and $C = 30$.

Finally, we examine the effect of varying $L_{\max}$ while fixing $L_{\min} = 20\%$. Figure 8 shows that expanding the eligibility range to include moderately charged bikes (up to 50–60%) improves profit and objective, as it broadens the candidate pool without excessive detours. Beyond this point, gains plateau and additional swaps provide little benefit. Larger capacities once again capture more profit

at the cost of longer routes. Table 5 summarizes the sensitivity analysis for SOC threshold policies.

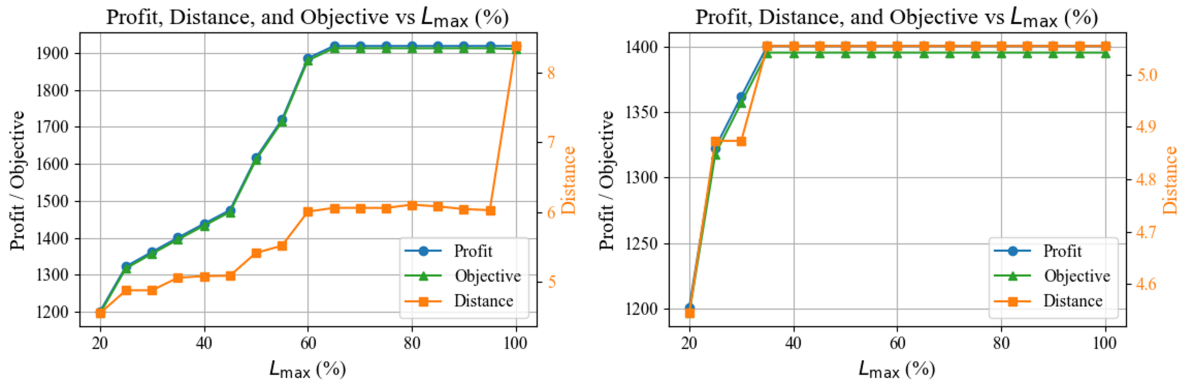

Figure 8: Sensitivity to $L_{\max}$ with $L_{\min} = 20\%$ for $C = 15$ and $C = 30$.

From a managerial standpoint, these findings suggest that threshold calibration should depend on fleet conditions and demand patterns. In denser service zones, where bikes are geographically clustered, a lower $L_{\min}$ can suffice because travel overhead is minimal. Conversely, in sparse or hilly areas, enforcing a higher $L_{\min}$ reduces the risk of running out of charged bikes before the next service cycle. The plateauing of profit at high $L_{\max}$ values further indicates a saturation effect: once moderately charged bikes are included, additional swaps no longer increase revenue proportionally, as the marginal gain per kilometer traveled declines.

Taken together, these results highlight the value of a balanced dual-threshold policy. The lower threshold ensures that the van prioritizes critically low-SOC bikes, while the upper threshold prevents inefficient interventions on bikes with high charge. Benchmark comparisons confirm that dual thresholds reduce unnecessary swaps and travel distance, with minimal sacrifice in profit. Sensitivity patterns further suggest that optimal thresholds lie around $L_{\min} = 40\%$ and $L_{\max} = 50$–60%, though precise tuning may depend on fleet size, capacity, and local demand. For operators, these findings imply that threshold design is not a purely technical parameter but a lever for balancing profitability, service reliability, and operational workload in real-world deployments.

The results demonstrate that carefully chosen SOC thresholds can yield substantial operational savings without sacrificing profit. The dual-threshold design supports selective interventions that

Table 5: Summary of sensitivity analysis results for SOC threshold policies

Policy	**Threshold Range Tested**	**Approximate Optimal Value(s)**	**Observed Effect on Profit and Distance**
Single-threshold (B1–B2)	$L_{\min} = 20$–30%	No clear optimum	High $L_{\min}$ caps profit potential, while low $L_{\min}$ leaves the van under-utilized; performance is sensitive to demand density.
Allow policy (Fig. 6)	$L_{\min} = 0$–80%	Stable until $L_{\min} \approx$ 40–50%	Profit and distance remain flat until mid-range, then decline sharply as fewer bikes qualify; larger capacity sustains profit but increases route length.
Force policy (Fig. 7)	$L_{\min} = 0$–80%	$L_{\min} \approx 40$–50%	Profit increases with stricter $L_{\min}$ as low-SOC bikes are prioritized, then drops once the eligible pool shrinks; larger capacity shifts the optimum slightly upward.
Dual-threshold (Original1–3)	$L_{\min} = 20\%$, $L_{\max} = 20$–100%	$L_{\max} \approx 50$–60%	Introducing $L_{\max}$ prevents unnecessary swaps on high-SOC batteries, improving efficiency with minor profit loss; benefits plateau beyond 60%.

maintain user availability while minimizing unnecessary travel. For example, the observed 5–10% reduction in total distance when applying $L_{\max}$ corresponds to meaningful cost savings in labor hours and fuel consumption.

In practical deployment, threshold policies could also adapt dynamically to daily demand or weather conditions. For instance, during peak commuting hours, lowering $L_{\min}$ slightly may prevent shortages in high-demand stations, while off-peak periods may favor stricter thresholds to conserve resources. Integrating these adaptive policies with predictive demand modeling could further enhance system responsiveness.

Finally, these findings reinforce that operational thresholds are not fixed constants but control levers. They allow operators to navigate trade-offs between profitability, sustainability, and battery longevity, each of which may be prioritized differently depending on local conditions or fleet

scale.

5.2 Sensitivity Analysis on Clustering Weight

The clustering weight parameter, denoted by α, controls the trade-off between spatial compactness and operational balance in the clustering refinement stage. As introduced in Section 4, lower values of α emphasize minimizing the geographic spread of bikes within clusters, while higher values emphasize balancing the distribution of low-SOC bikes across clusters. Since clustering directly affects the efficiency of subsequent routing optimization, the choice of α indirectly influences profit, travel distance, and the number of feasible swaps.

To examine this effect, the clustering–routing workflow was executed across a range of α values. Figure 9 illustrates the relationship between α and the total distance from bikes to their assigned cluster centers. The curve exhibits a U-shaped pattern, with a distinct minimum at $\alpha = 0.6$. This value represents the most effective trade-off between spatial cohesion and operational balance. The observed U-shaped relationship highlights a fundamental trade-off in spatial optimization: compact clusters minimize intra-cluster travel but risk overloading certain service regions, while balanced clusters spread demand evenly at the cost of longer routes. The parameter α therefore serves as a tuning mechanism for prioritizing geographical versus operational balance. In a practical setting, operators in dense downtown areas may benefit from a slightly lower α to exploit proximity and reduce routing time, while suburban or less dense systems may prefer higher α values to ensure even service coverage across a wider area.

At very low α values, clusters are geographically compact, but they often contain disproportionate numbers of low-SOC bikes. This imbalance burdens some clusters with excessive demand while leaving others under-utilized, reducing overall profitability and creating inefficiencies in route assignment. Conversely, very high α values yield balanced clusters in terms of SOC distribution, but they sacrifice spatial compactness. In this regime, routes become longer and less efficient because e-bikes assigned to the same cluster are geographically dispersed. Table 6 summarizes the effect of the weight parameter on the system performance.

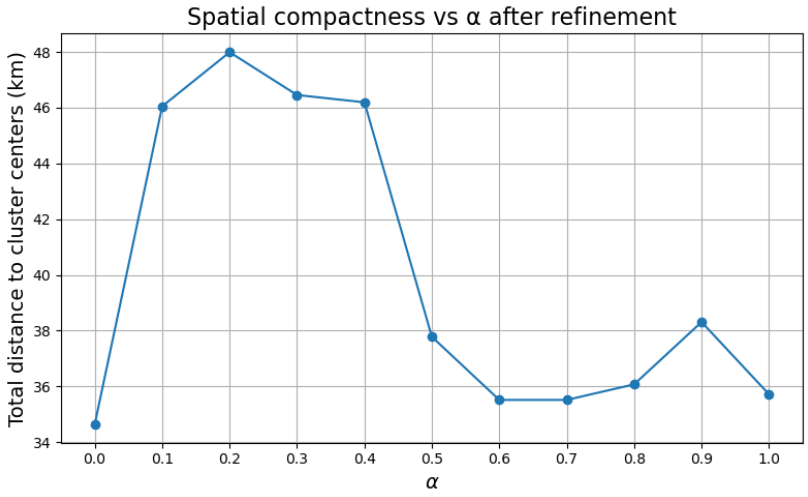

Figure 9: Spatial compactness versus α after refinement. The y-axis measures the total distance (in km) from all e-bikes to their assigned cluster centers. Lower values indicate more compact clusters.

The optimal weight of $\alpha = 0.6$ identified in this analysis was adopted for all baseline experiments reported in Sections 5–6. This ensures that the clustering stage provides a balanced foundation for route optimization, avoiding extreme scenarios of purely compact or purely balanced clusters. From an operational standpoint, this finding suggests that fleet managers benefit from considering both spatial and SOC dimensions simultaneously: clustering on geography alone can overload routes with low-battery bikes, while clustering purely on SOC balance can create impractically long tours. A moderate weighting captures the advantages of both approaches and supports robust system performance.

The choice of $\alpha = 0.6$ not only minimizes intra-cluster distance but also stabilizes route performance across different SOC distributions. This parameter was subsequently fixed for all further experiments to ensure comparability across case studies. Moreover, the robustness of this value suggests that the proposed clustering refinement method can generalize well to cities with different spatial densities, as long as α is tuned within a moderate range around the optimal point.

5.3 Summary of Sensitivity Findings

Across all tests, the sensitivity analysis reveals that system performance is most influenced by three parameters: the lower SOC threshold $L_{\min}$, van capacity C, and the clustering weight α. Adjusting

Table 6: Summary of sensitivity analysis results for clustering weight parameter α

Parameter	**Range Tested**	**Optimal Value**	**Observed Effect on System Performance**
Clustering weight α	0–1 (tested in increments of 0.1)	$\alpha \approx 0.6$	• Low α: clusters are compact geographically but imbalanced in SOC, overloading some routes. • High α: clusters are balanced in SOC but geographically dispersed, increasing travel distance. • Mid-range α: balances compactness and SOC distribution, minimizing route distance while keeping workloads feasible.

these parameters reshapes both routing structure and profitability. The results suggest that:

- Optimal thresholds ($L_{\min} = 40$–50%, $L_{\max} = 50$–60%) maintain high service reliability while minimizing redundant trips.
- Increasing van capacity yields higher profit but with diminishing returns once C exceeds regional demand.
- The clustering weight $\alpha = 0.6$ provides the best trade-off between spatial cohesion and workload balance.

Collectively, these findings demonstrate that the proposed framework is robust to moderate parameter variation and that fine-tuning these key parameters can enhance operational efficiency without requiring structural changes to the optimization model.

6 Case Study: San Francisco

To demonstrate the practical application of the proposed swapping–routing framework, this chapter evaluates the model using data from the Bay Wheels dockless e-bike system in San Francisco, California. The city provides an ideal and well-documented test environment for micromobility research, having been one of the earliest and most closely studied U.S. cities in the large-scale adoption of shared e-bikes and e-scooters [28, 31]. Its dense urban structure, steep elevation gradients, and uneven trip distribution make battery management a particularly complex operational challenge. These characteristics allow for a rigorous assessment of the framework's ability to handle spatial heterogeneity, constrained vehicle routing, and dynamic energy conditions.

San Francisco's topography has long been recognized as one of the most defining and operationally challenging features of its micromobility network. Steep hills in neighborhoods such as Nob Hill and Twin Peaks cause highly asymmetric energy consumption, where bikes traveling uphill deplete batteries rapidly while those in flatter or downhill regions maintain higher states of charge (SOC). This pattern creates a spatially imbalanced distribution of charge levels across the city, requiring operators to strategically allocate swapping resources to maintain service reliability. Furthermore, trip demand fluctuates substantially across time and space, with heavy usage in the downtown and waterfront districts and lighter activity in residential and peripheral zones. These conditions collectively test the framework's capacity to generate adaptive policies that remain both efficient and interpretable.

Previous research on San Francisco's micromobility ecosystem has primarily focused on behavioral, regulatory, and environmental aspects rather than operational optimization. For example, Mooney et al. [31] analyzed the deployment and governance of dockless electric scooters, highlighting the regulatory challenges of managing new modes of shared mobility in dense urban settings. Similarly, Lazarus et al. [28] examined how shared micromobility services, including e-bikes and scooters, influenced user travel choices and contributed to shifts in urban transportation patterns. While these studies provide valuable insights into user behavior and policy frameworks, they do not address the underlying logistical questions of fleet balancing and energy replenishment.

Building upon this foundation, the present study extends the analysis to the operator's perspective, focusing on how vehicle-level optimization can improve the efficiency and sustainability of shared e-bike systems. By integrating clustering, routing, and threshold refinement into a unified model, this case study explores how strategic control of SOC-based swap policies can yield stable and interpretable operational outcomes. San Francisco's open GBFS data feeds, which provide detailed real-time and historical information on bike status and charge levels, offer an exceptional empirical basis for this evaluation.

In summary, San Francisco combines the necessary data transparency, physical complexity, and operational diversity to serve as a representative and demanding environment for testing the proposed battery swapping and routing framework. Its mixture of dense urban corridors, steep hills, and high ridership density enables the assessment of both global threshold behavior and localized decision-making under realistic service conditions.

The data used in this study were collected exclusively from the Bay Wheels GBFS live feeds, which provide real-time bike-level information including unique bike identifiers, geographic coordinates, state-of-charge (SOC), and availability status. The feed is updated approximately every five minutes, offering a continuously refreshed view of fleet conditions across the city.

Before analysis, all records were cleaned to ensure internal consistency. Bikes with missing or duplicate timestamps, implausible coordinates (e.g., locations outside city limits), or abnormal SOC readings (less than 0% or greater than 100%) were removed. Coordinates were projected from WGS84 latitude–longitude to UTM Zone 10N to permit distance calculations in meters, and geodesic distances were used for bike-to-bike and bike-to-depot separation to minimize distortion across the city's irregular surface.

Because operational optimization requires a static spatial snapshot, the live feed was aggregated into a single morning snapshot representing fleet conditions at the start of a typical daily service cycle. This snapshot serves as the input to both the clustering and routing stages. The resulting dataset contains 280 operational e-bikes distributed across San Francisco's service area, spanning both dense downtown corridors and lower-density peripheral neighborhoods.

The first stage of the framework partitions the fleet into geographically compact service regions to enable localized routing and decision-making. This is achieved through a capacity-constrained k-means algorithm, which groups the 280 e-bikes into ten clusters ($k = 10$) while enforcing a per-cluster capacity of 35 bikes as can be seen in Figure 11. This constraint mirrors the workload a single service van can realistically manage, ensuring that the optimization at the routing stage remains computationally feasible.

The k-means algorithm minimizes within-cluster squared distances subject to capacity constraints, balancing geographic compactness and workload distribution. However, because initial centroid assignments may not perfectly reflect operational realities such as uneven terrain or local demand, an additional refinement stage is implemented. In this step, single-bike relocations are evaluated iteratively: for each bike, the model tests whether moving it to a neighboring cluster within a 5 km radius improves both its local objective score (based on SOC deviation and proximity) and the overall sum of cluster scores. If both conditions are met, the move is accepted.

This relocation heuristic allows clusters to adapt to local imbalances, especially those caused by uneven energy depletion or topographic separation. For instance, bikes located near steep gradients are often reassigned to clusters whose centroids better represent their physical accessibility for the service vehicle. The weight parameter α in the combined objective (balancing distance and SOC deviation) is fixed at 0.60, following the results of the sensitivity analysis in Section 5.2.

After refinement, each cluster represents a compact, topographically consistent service area with roughly balanced SOC distributions, forming the input for routing optimization.

Within each refined cluster, routing and swapping decisions are jointly optimized using a Mixed-Integer Linear Programming (MILP) model that incorporates Miller–Tucker–Zemlin (MTZ) sub-tour elimination constraints. The optimization seeks to maximize profit while minimizing travel cost, formalized by the objective function in Equation 2, which combines revenue from battery swaps and distance-dependent operating cost.

The MILP is solved independently for each cluster using a Python-based optimization framework implemented with the PuLP library, on an Intel i7 processor with 16 GB of RAM. Average runtime

per cluster is approximately 1–2 minutes, depending on cluster size and threshold tightness.

Bikes with SOC below $L_{\min}$ are automatically included in the route as forced swaps, while those with SOC above $L_{\max}$ are excluded. Bikes with intermediate SOC values are evaluated endogenously by the optimizer, which decides whether servicing them increases total net benefit. This allows the MILP to adapt to economic trade-offs between battery availability, distance, and marginal gains from additional swaps.

The dual-threshold parameters $(L_{\min}, L_{\max})$ evolve through an outer iterative refinement loop. After solving the MILP for all clusters, the model computes average improvement in the objective function with respect to the current thresholds. The thresholds are then updated in 5-percentage-point increments following a steepest-descent approach—if tightening or relaxing the bounds yields higher profit, the model moves in that direction.

This iterative process continues until changes in global objective value fall below 1% between consecutive iterations, indicating convergence. The algorithm is initialized with quartile-based values $(L_{\min}, L_{\max}) = (10.73\%, 22.40\%)$ and executed for ten outer iterations.

In the first iteration, six one-bike relocations were accepted (bikes 254, 94, 215, 57, 179, and 161 moving from cluster 1 to clusters 4 or 7). Subsequent iterations tested additional relocations but accepted few, as most yielded negligible improvements. Over the course of refinement, thresholds increased steadily: $L_{\min}$ rose from 10.73% to 19.23%, and $L_{\max}$ from 22.40% to 40.40%. By iteration 10, both global thresholds and cluster assignments had stabilized, demonstrating convergence of the joint refinement–optimization loop.

The convergence trajectory (Figure 10) shows that most of the threshold adjustment occurs within the first three iterations, after which only minor corrections are made. This rapid stabilization highlights the efficiency of the refinement procedure, which requires only limited computational effort to reach equilibrium.

Cluster-level outcomes reveal substantial heterogeneity across San Francisco's geography. Northern and downtown clusters exhibited higher swap frequencies and larger objective values due to dense trip generation and steeper gradients. In contrast, flatter southern clusters with lower demand

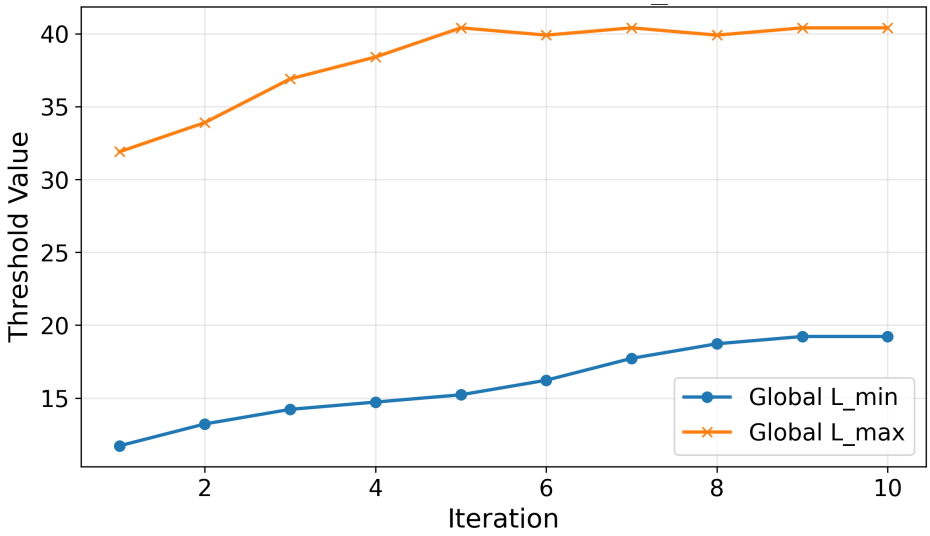

Figure 10: Convergence of global SOC thresholds in the San Francisco case study. Initialization begins from quartile-based values $(10.73\%, 22.40\%)$ and stabilizes after ten iterations at approximately $(19.23\%, 40.40\%)$.

required fewer swaps and achieved lower objective values.

Some clusters temporarily adopted very low $L_{\min}$ values (1–7%) with moderate $L_{\max}$ around 30–35%, while others stabilized with higher thresholds near $(19\%, 45\%)$. This pattern reflects adaptive behavior: in high-demand zones, the optimizer prioritizes aggressive replacement of depleted batteries, whereas in quieter regions, swaps are deferred until SOC falls below critical levels.

A few specific bikes, such as bike 17 in cluster 0 and bike 53 in cluster 9, repeatedly emerged as low-performing outliers during iterations. These cases indicate microregions where parking density or demand imbalance causes persistent inefficiency, valuable diagnostic information for operators planning depot adjustments or service zone redesigns.

Table 7 summarizes the final results after convergence. Each cluster's MILP solution typically filled the truck capacity (20 swaps), with 11–19 forced swaps per cluster and additional intermediate decisions chosen by the optimizer. The variation in objective values reflects both spatial demand differences and the marginal benefits of servicing each area.

To provide a system-level view of the final solution, Figure 12 shows the optimized routes for all ten clusters in the San Francisco case study on a single map. Each color represents one refined cluster, and every route starts and ends at the depot. The figure highlights how the capacity-constrained

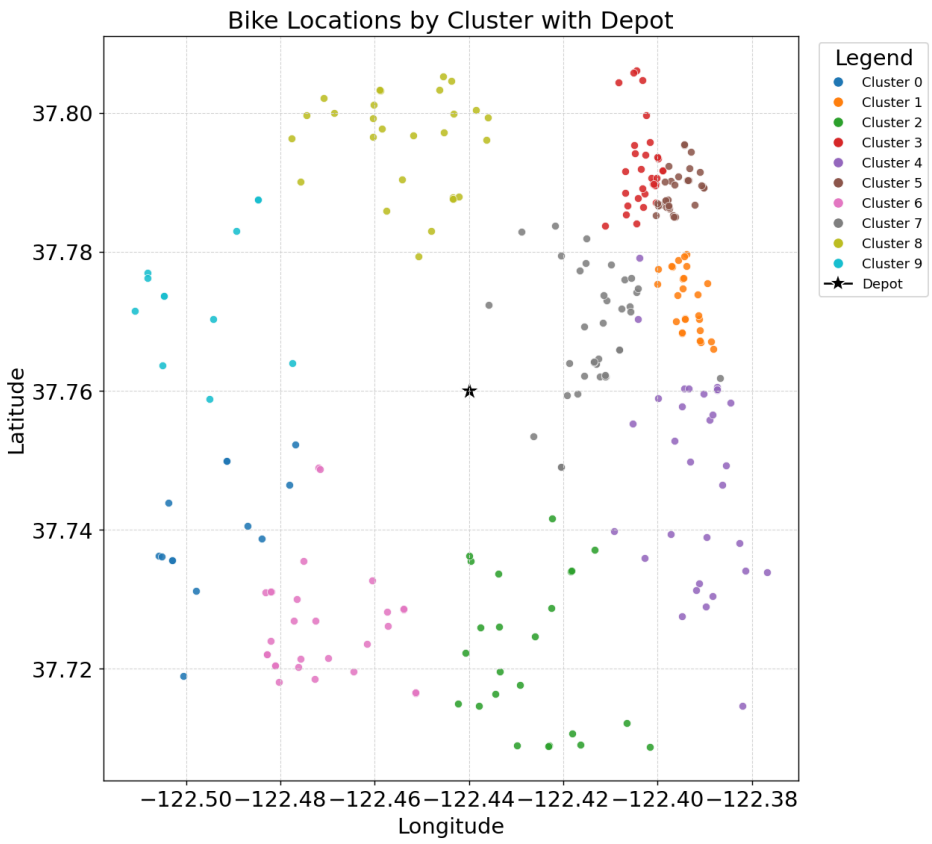

Figure 11: Spatial distribution of the ten refined service clusters in San Francisco. Clusters are geographically compact and align with natural travel corridors shaped by topography, demand, and accessibility.

clustering and refinement stages partition the fleet into compact service regions that can be routed independently. Routes associated with downtown and northern clusters appear shorter and denser, reflecting high bike density and more frequent low-SOC occurrences, while peripheral clusters cover larger areas with longer inter-stop links. Overall, the plot illustrates the main operational advantage of the proposed framework: a large city-scale swapping problem is decomposed into multiple localized tours that remain feasible under van capacity limits and can be executed in parallel.

To illustrate the spatial outcomes of the optimization process, Figures 13 and 14 display the optimized service routes for all ten clusters in San Francisco. Each subplot shows the bikes belonging to a given cluster, the optimized van route, and the subset of bikes selected for swapping (highlighted with orange squares). The titles of each subplot report the cluster identifier, lower and upper SOC thresholds $(L_{\min}, L_{\max})$, and the final objective value of the corresponding MILP solution. These

Table 7: Cluster-level operational outcomes and objective values in the San Francisco case study

Cluster	Bikes	Forced Swaps	Forced No-Swaps	MILP Decisions	MILP Swapped	Obj. Value
0	13	3	0	10	13	1673.3
1	29	16	1	12	20	3224.4
2	24	6	1	17	20	2884.4
3	35	19	3	13	20	3603.0
4	33	17	0	16	20	3465.1
5	20	8	0	12	20	3349.5
6	23	10	2	11	20	3255.6
7	25	11	1	13	20	3354.7
8	24	10	0	14	20	3437.1
9	21	7	0	14	20	3215.6

visualizations help reveal how the routing model adapts to spatial heterogeneity and SOC-based decisions.

Each subplot corresponds to one cluster, showing the optimized van route and the swapped bikes (orange squares).

Overall, the route visualizations demonstrate that the clustering and routing components jointly produce compact, non-overlapping service regions that correspond well to San Francisco's urban structure. Clusters located near the downtown core and northern hills (Clusters 1–4) show dense bike concentrations and relatively short inter-stop distances, while those located toward the southern and coastal regions (Clusters 8–9) cover larger spatial extents with longer travel segments between stops. This spatial variation directly reflects the city's topography and demand distribution: higher ridership areas generate more frequent swaps, whereas flatter or lower-demand zones require fewer interventions.

The first three clusters, located in the northeastern portion of the city, represent high-demand zones with numerous low-SOC bikes. Their optimized routes are compact, featuring many short links between adjacent stops. These routes efficiently capture localized replacement needs while maintaining minimal travel redundancy. Objective values for these clusters are moderate (ranging between 1600–2900), balancing travel cost and revenue from frequent swaps.

Clusters 3 and 4 lie near the central and waterfront districts, where trip density is high and terrain

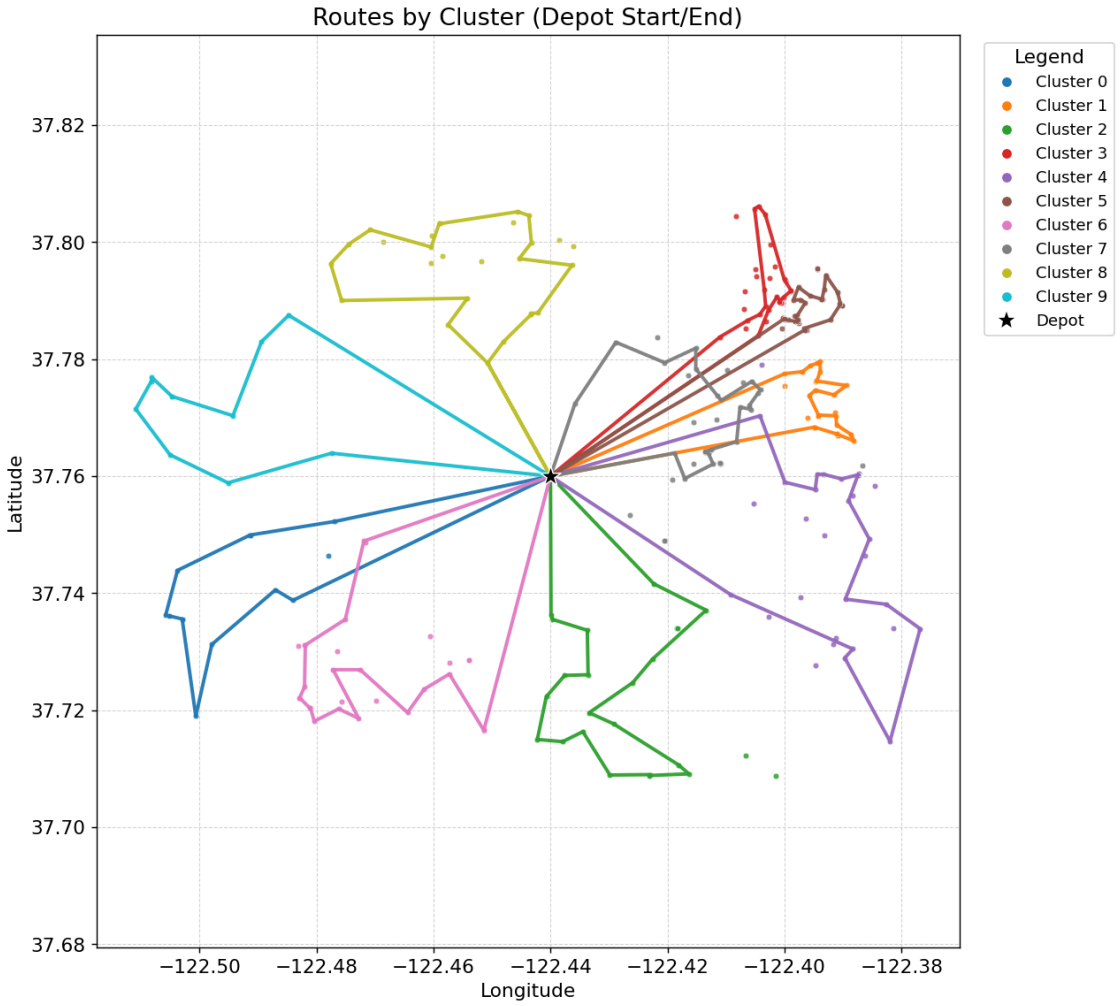

Figure 12: Optimized van routes for all refined clusters in the San Francisco case study. Each color shows one cluster tour, with all routes starting and ending at the depot. The visualization summarizes the final spatial structure of the solution before presenting cluster-level route details.

is moderately steep. The optimized paths show tightly packed routes with multiple small loops, consistent with regions experiencing sustained battery depletion due to uphill usage. These clusters exhibit some of the highest objective values (around 3400–3600), indicating strong operational payoff for servicing these areas.

These mid-city clusters exhibit more elongated routes that stretch across adjacent neighborhoods. The presence of a few longer connecting links suggests that the algorithm balances local compactness with the need to reach scattered low-SOC bikes. Cluster 7, for instance, covers a relatively large region but still maintains a coherent route structure, yielding an objective of 3935.4—the

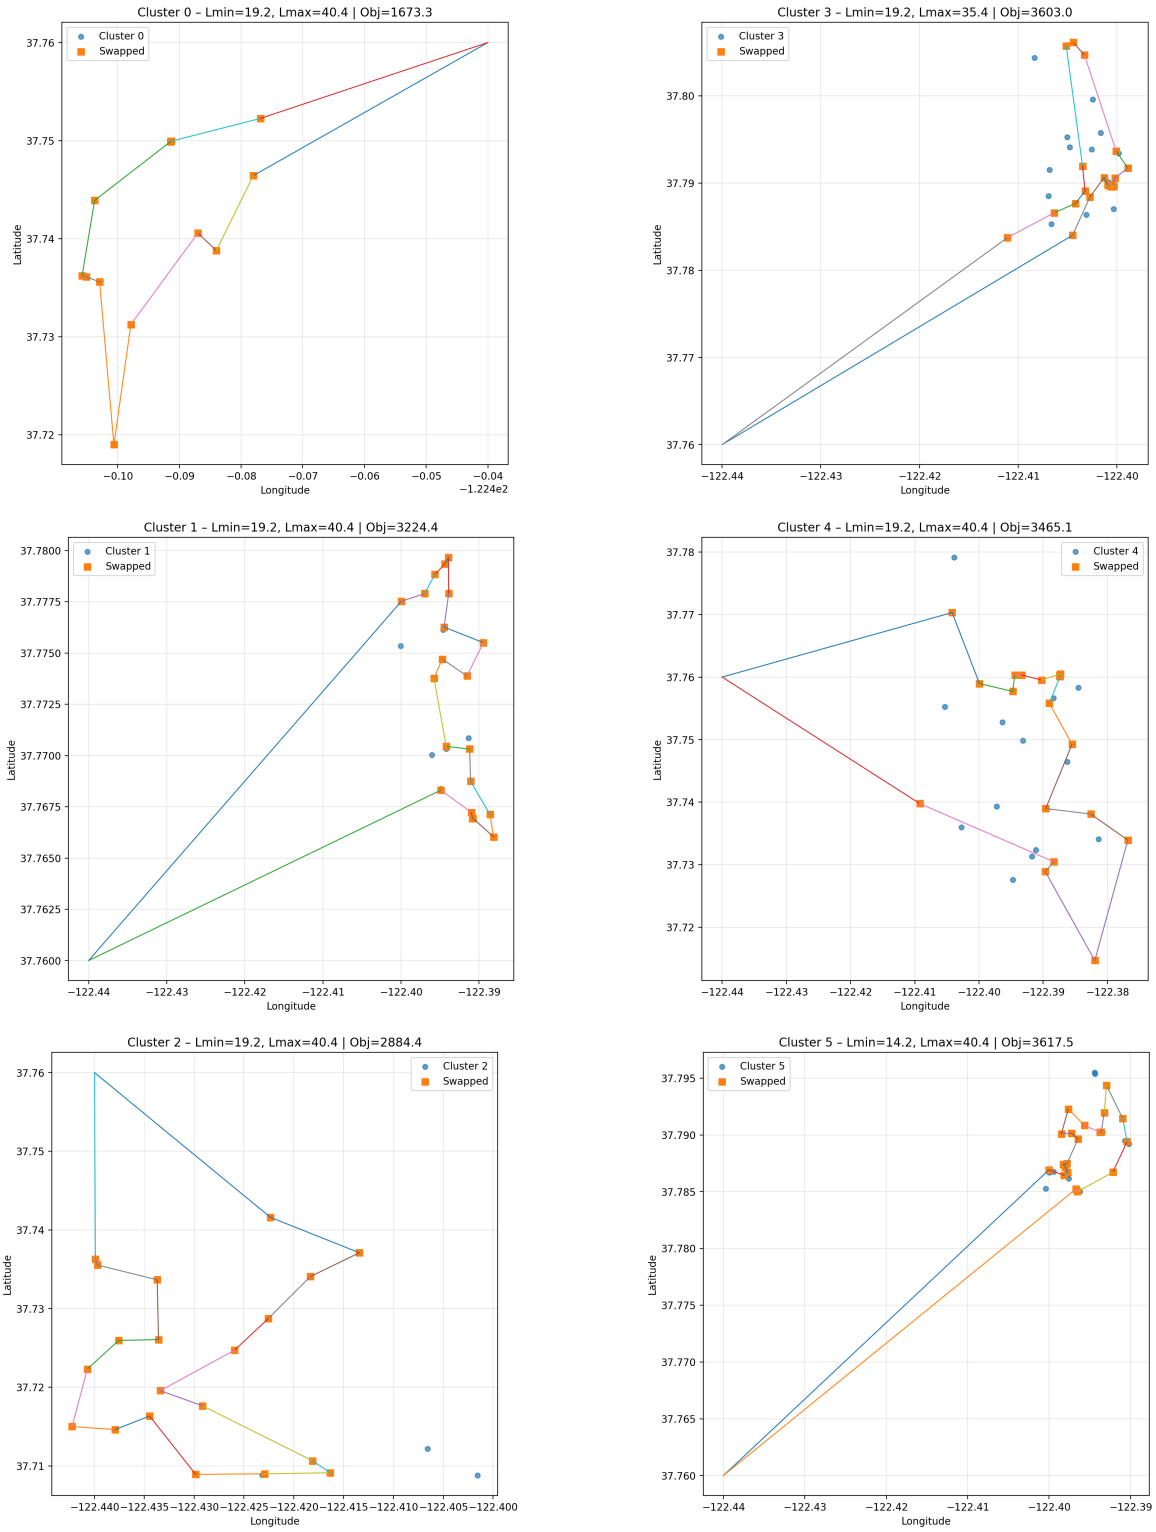

Figure 13: Optimized routes for Clusters 0–5 in San Francisco. Each panel shows the optimized van path and swapped bikes.

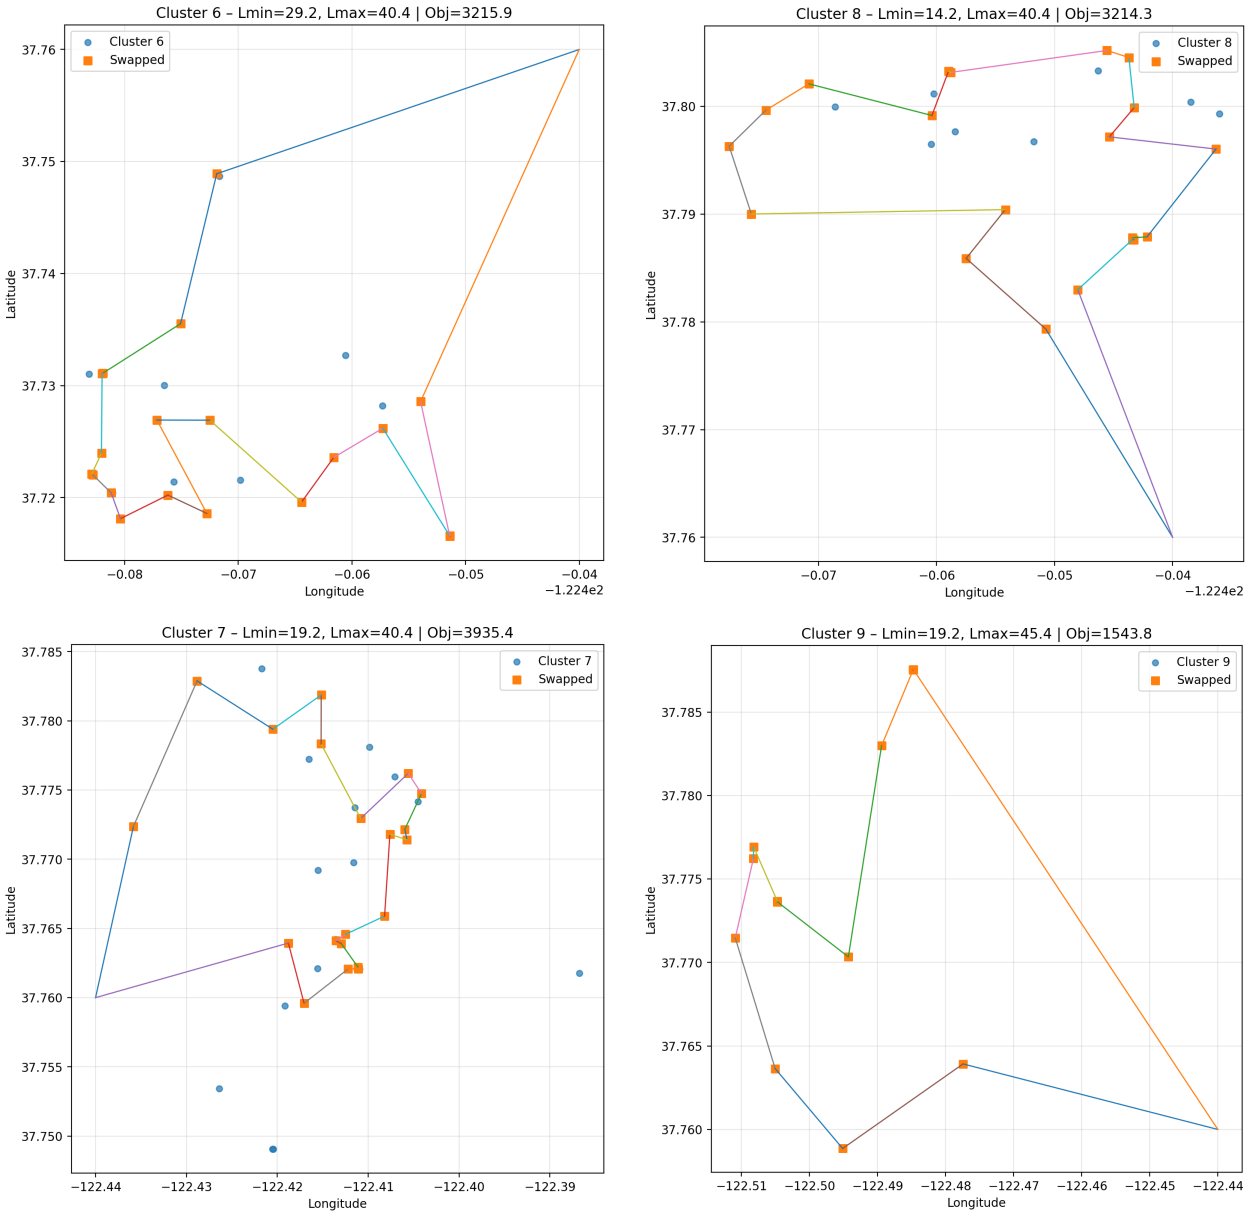

Figure 14: Optimized service routes: two columns of two (Clusters 6–7 left, 8–9 right).

highest among all clusters. This reflects the economic efficiency of serving a mixed-density zone with several high-priority swaps.

The final two clusters correspond to the southern and western periphery of the service area, characterized by lower trip density and flatter topography. Their routes span wider geographic areas with fewer stops, and the number of forced swaps is smaller than in the northern clusters. The lower objective values (around 1500–3200) are consistent with reduced operational activity and longer average distances between bikes. Nonetheless, both clusters remain geographically cohesive, confirming that the refinement procedure avoided spatial fragmentation.

Across all clusters, routes tend to follow natural travel corridors and avoid unnecessary detours, indicating that the MILP formulation successfully captures distance-based efficiency. The number of bikes selected for swapping nearly matches the 20-unit van capacity in most clusters, showing that the model uses available resources fully while respecting SOC thresholds. Variations in $(L_{\min}, L_{\max})$ among clusters reflect local operational needs: downtown regions adopt higher upper limits to prevent over-servicing, while peripheral regions tolerate slightly lower lower limits to maintain adequate coverage.

From a managerial standpoint, these visual results confirm that the framework generates routes that are both interpretable and actionable. Each map can be viewed as a service plan for a specific van, guiding technicians through a geographically consistent route that maximizes the operational benefit. The figures also highlight potential micro-regions where repeated low-SOC occurrences suggest infrastructure or demand imbalances, offering valuable diagnostic information for future system design.

The results demonstrate that a global swapping policy with $L_{\min} \approx 13\%$ and $L_{\max} \approx 38$–40% is practical and effective for San Francisco's operational conditions. The model provides both a global decision rule and local flexibility through MILP optimization, allowing operators to interpret and apply thresholds easily in the field.

Capacity-aware clustering and single-bike refinement ensure that clusters remain geographically compact, improving van routing efficiency and reducing redundant travel. The algorithm's conver-

gence after only ten iterations indicates that operators could deploy this model in an iterative, near real-time planning environment with modest computational effort.

From a strategic perspective, the framework offers more than just optimized routes—it delivers diagnostic insights. Clusters or bikes that repeatedly trigger low performance can be flagged for deeper investigation, possibly indicating misaligned parking zones, inadequate charging coverage, or spatial mismatches between supply and demand. The transparency of this iterative process also supports reproducibility and auditability, allowing planners to trace how threshold adjustments influence fleet outcomes.

7 Limitations and Future Work

While the San Francisco case study demonstrates the stability, convergence, and interpretability of the proposed framework, several limitations remain. First, the analysis is based on static snapshots of state-of-charge (SOC) and fleet conditions. In practice, SOC evolves continuously as users interact with the system during service operations, introducing temporal coupling between demand, depletion, and swapping decisions.

Second, the routing component relies on mixed-integer linear programming solved under fixed computational time limits. Although this ensures tractability for large instances, it may prevent the solver from reaching global optimality in dense urban settings. Third, the operational setting assumes a single service van per cluster, whereas real-world fleet management typically involves multiple vehicles operating simultaneously, often sharing depots and overlapping service areas.

Finally, several economic parameters, including labour cost, per-swap revenue, and battery degradation effects, were simplified or calibrated using representative rather than proprietary data. These assumptions were adopted to emphasize methodological clarity but may limit direct cost accuracy in operational deployment.

These limitations define several promising directions for future research. A key extension is the incorporation of time-dependent SOC evolution and stochastic demand into the clustering–routing loop. Embedding the framework within a rolling-horizon or event-driven simulation would allow policies to adapt dynamically to intra-day variations and real-time system states.

Another important direction is the extension to multi-vehicle routing. Allowing multiple service vans to operate concurrently, either cooperatively or competitively, would significantly improve operational realism and scalability. This could include shared depots, coordinated scheduling, or dynamic reassignment of clusters as system conditions evolve.

Future work could also integrate decisions related to fixed charging infrastructure, enabling joint optimization between mobile battery swapping and stationary charging investments. Such integration would provide a more holistic view of energy replenishment strategies in micromobility systems.

From a computational perspective, advanced solution techniques such as decomposition methods, metaheuristics, or learning-based approaches could be explored to improve scalability and reduce solution times. These methods may enable near real-time optimization in large metropolitan networks where exact MILP solutions become computationally prohibitive.

Finally, expanding the framework to multiple cities would allow comparative analysis across different urban contexts. Differences in geography, topography, trip demand, and fleet composition could be leveraged to derive generalizable insights into how local conditions shape optimal battery swapping and routing policies.

8 Conclusion

This thesis presented an integrated optimization framework for battery swapping in dockless e-bike sharing systems, addressing the dual challenge of operational efficiency and profitability in large-scale micromobility networks. The proposed methodology combines capacity-constrained clustering with mixed-integer linear programming (MILP)-based routing within an iterative, feedback-driven loop, linking spatial service design with routing and swapping decisions.

The results demonstrate that this combined clustering–routing approach improves coordination between spatial balance and routing efficiency. The introduction of a dual-threshold SOC policy proved superior to traditional single-threshold strategies by reducing unnecessary swaps while maintaining profitability. Sensitivity analysis further showed that balanced weighting between spatial compactness and SOC uniformity yields the most efficient system performance. The San Francisco case study confirmed that the framework is robust and convergent under real-world conditions, producing stable cluster configurations, interpretable SOC thresholds, and efficient routing solutions.

The main contributions of this thesis are summarized as follows:

- Developed a two-stage iterative optimization framework that integrates capacity-constrained clustering with routing optimization for battery swapping in dockless e-bike systems.
- Introduced and validated a dual-threshold SOC policy that balances operational efficiency, profitability, and service quality more effectively than single-threshold approaches.
- Designed a cluster scoring and refinement mechanism that jointly accounts for spatial compactness and SOC heterogeneity, and demonstrated its effectiveness through sensitivity analysis.
- Validated the framework using real operational data from San Francisco's Bay Wheels system, yielding interpretable and practically actionable insights for fleet operators.

Overall, this research demonstrates that integrating clustering refinement with optimization-based

routing provides a robust and interpretable approach to managing battery swapping operations in dockless e-bike systems. By explicitly capturing spatial, operational, and economic interdependencies, the framework bridges theoretical optimization and practical implementation. The insights gained from this work support data-driven decision-making for micromobility operators and establish a strong foundation for future advances in battery management and shared electric mobility systems.

Bibliography

[1] National Association of City Transportation Officials, "Shared micromobility in 2023," 2024. [Online]. Available: https://nacto.org/publication/shared-micromobility-in-2023/

[2] A. Fyhri and N. Fearnley, "Effects of e-bikes on bicycle use and mode share," *Transportation research part D: transport and environment*, vol. 36, pp. 45–52, 2015.

[3] N. Popovich, E. Gordon, Z. Shao, Y. Xing, Y. Wang, and S. Handy, "Experiences of electric bicycle users in the sacramento, california area," *Travel Behaviour and Society*, vol. 1, no. 2, pp. 37–44, 2014.

[4] A. Bassolas, J. Grau-Escolano, and J. Vicens, "Spatiotemporal variability and prediction of e-bike battery levels in bike-sharing systems," *Scientific Reports*, vol. 15, no. 1, p. 7171, 2025.

[5] H. Wu, "A survey of battery swapping stations for electric vehicles: Operation modes and decision scenarios," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 8, pp. 10 163–10 185, 2021.

[6] S. Shao, D. Li, Y. Zhou, and J.-B. Sheu, "Optimizing battery swapping for city-scale e-bike sharing systems: A three-stage spatial–temporal cluster-based approach," *Transportation Research Part E: Logistics and Transportation Review*, vol. 198, p. 104145, 2025.

[7] M. Xu, Y. Di, Z. Zhu, H. Yang, and X. Chen, "Designing van-based mobile battery swapping and rebalancing services for dockless ebike-sharing systems based on the dueling double deep q-network," *Transportation Research Part C: Emerging Technologies*, vol. 138, p. 103620, 2022.

[8] X. Xie, X. Dai, and Z. Pei, "Empowering the capillary of the urban daily commute: Battery deployment analysis for the locker-based e-bike battery swapping," *Transportation Science*, vol. 58, no. 1, pp. 176–197, 2024.

[9] Y. Zhou, Z. Lin, R. Guan, and J.-B. Sheu, "Dynamic battery swapping and rebalancing strategies for e-bike sharing systems," *Transportation Research Part B: Methodological*, vol. 177, p. 102820, 2023.

[10] P. DeMaio, "Bike-sharing: History, impacts, models of provision, and future," *Journal of public transportation*, vol. 12, no. 4, pp. 41–56, 2009.

[11] S. A. Shaheen, S. Guzman, and H. Zhang, "Bikesharing in europe, the americas, and asia: past, present, and future," *Transportation research record*, vol. 2143, no. 1, pp. 159–167, 2010.

[12] T. Mátrai and J. Tóth, "Comparative assessment of public bike sharing systems," *Transportation research procedia*, vol. 14, pp. 2344–2351, 2016.

[13] S. Shaheen, A. Cohen, and J. Broader, "What's the'big'deal with shared micromobility? evolution, curb policy, and potential developments in north america," *Built Environment*, vol. 47, no. 4, pp. 499–514, 2021.

[14] J. Pucher and R. Buehler, "Cycling for everyone: lessons from europe," *Transportation research record*, vol. 2074, no. 1, pp. 58–65, 2008.

[15] E. Fishman, "Bikeshare: A review of recent literature," *Transport reviews*, vol. 36, no. 1, pp. 92–113, 2016.

[16] Y. Sun, "Sharing and riding: How the dockless bike sharing scheme in china shapes the city," *Urban Science*, vol. 2, no. 3, p. 68, 2018.

[17] S. Shaheen and N. Chan, "Mobility and the sharing economy: Potential to facilitate the first- and last-mile public transit connections," *Built environment*, vol. 42, no. 4, pp. 573–588, 2016.

[18] X. Ma, X. Zhang, X. Li, X. Wang, and X. Zhao, "Impacts of free-floating bikesharing system on public transit ridership," *Transportation research part D: transport and environment*, vol. 76, pp. 100–110, 2019.

[19] S. Jäppinen, T. Toivonen, and M. Salonen, "Modelling the potential effect of shared bicycles on public transport travel times in greater helsinki: An open data approach," *Applied geography*, vol. 43, pp. 13–24, 2013.

[20] B. Caulfield, M. O'Mahony, W. Brazil, and P. Weldon, "Examining usage patterns of a bike-sharing scheme in a medium sized city," *Transportation research part A: policy and practice*, vol. 100, pp. 152–161, 2017.

[21] NACTO, "Shared micromobility report: 2020-2021," 2022. [Online]. Available: https://nacto.org/publication/shared-micromobility-report-2020-2021/

[22] R. L. Abduljabbar, S. Liyanage, and H. Dia, "The role of micro-mobility in shaping sustainable cities: A systematic literature review," *Transportation research part D: transport and environment*, vol. 92, p. 102734, 2021.

[23] J. F. Teixeira and M. Lopes, "The link between bike sharing and subway use during the covid-19 pandemic: The case-study of new york's citi bike," *Transportation research interdisciplinary perspectives*, vol. 6, p. 100166, 2020.

[24] A. Nikitas and E. Bakogiannis, "Urban transport, resilient cities and covid-19: Testing mobility interventions for a disrupted world," p. 105183, 2024.

[25] C. Cherry and R. Cervero, "Use characteristics and mode choice behavior of electric bike users in china," *Transport policy*, vol. 14, no. 3, pp. 247–257, 2007.

[26] A. Bigazzi and K. Wong, "Electric bicycle mode substitution for driving, public transit, conventional cycling, and walking," *Transportation research part D: transport and environment*, vol. 85, p. 102412, 2020.

[27] H. Luo, Z. Kou, F. Zhao, and H. Cai, "Comparative life cycle assessment of station-based and dock-less bike sharing systems," *Resources, Conservation and Recycling*, vol. 146, pp. 180–189, 2019.

[28] J. Lazarus, J. C. Pourquier, F. Feng, H. Hammel, and S. Shaheen, "Micromobility evolution and expansion: Understanding how docked and dockless bikesharing models complement and compete–a case study of san francisco," *Journal of Transport Geography*, vol. 84, p. 102620, 2020.

[29] Y. Zhang and Z. Mi, "Environmental benefits of bike sharing: A big data-based analysis," *Applied energy*, vol. 220, pp. 296–301, 2018.

[30] E. Fishman and P. Schepers, "Global bikeshare: what the data tells us about safety," in *International Cycling Safety Conference (ICSC2014), 3rd, 2014, Gothenburg, Sweden*, 2014.

[31] S. J. Mooney, K. Hosford, B. Howe, A. Yan, M. Winters, A. Bassok, and J. A. Hirsch, "Freedom from the station: Spatial equity in access to dockless bike share," *Journal of transport geography*, vol. 74, pp. 91–96, 2019.

[32] D. V. Nixon and T. Schwanen, "Bike sharing beyond the norm," *Journal of Transport Geography*, vol. 80, p. 102492, 2019.

[33] Y. Zhang, D. Kasraian, and P. van Wesemael, "Built environment and micro-mobility," *Journal of Transport and Land Use*, vol. 16, no. 1, pp. 293–317, 2023.

[34] E. Fishman, S. Washington, and N. Haworth, "Bikeshare's impact on active travel: Evidence from the united states, great britain, and australia," *Journal of Transport & Health*, vol. 2, no. 2, pp. 135–142, 2015.

[35] E. Fishman, P. Schepers, and C. B. M. Kamphuis, "Dutch cycling: quantifying the health and related economic benefits," *American journal of public health*, vol. 105, no. 8, pp. e13–e15, 2015.

[36] J. Woodcock, M. Tainio, J. Cheshire, O. O'Brien, and A. Goodman, "Health effects of the london bicycle sharing system: health impact modelling study," *Bmj*, vol. 348, 2014.

[37] I. Otero, M. J. Nieuwenhuijsen, and D. Rojas-Rueda, "Health impacts of bike sharing systems in europe," *Environment international*, vol. 115, pp. 387–394, 2018.

[38] S. A. Shaheen and A. P. Cohen, "Carsharing and personal vehicle services: worldwide market developments and emerging trends," *International journal of sustainable transportation*, vol. 7, no. 1, pp. 5–34, 2013.

[39] S. Howland, N. McNeil, J. Broach, K. Rankins, J. MacArthur, and J. Dill, "Breaking barriers to bike share: Insights on equity from a survey of bike share system owners and operators," 2017.

[40] E. Desjardins, C. D. Higgins, and A. Páez, "Examining equity in accessibility to bike share: A balanced floating catchment area approach," *Transportation research part D: transport and environment*, vol. 102, p. 103091, 2022.

[41] X. Qian, M. Jaller, and G. Circella, "Equitable distribution of bikeshare stations: An optimization approach," *Journal of transport geography*, vol. 98, p. 103174, 2022.

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

[43] C. M. De Chardon, G. Caruso, and I. Thomas, "Bike-share rebalancing strategies, patterns, and purpose," *Journal of transport geography*, vol. 55, pp. 22–39, 2016.

[44] J. Zhang and M. Meng, "Bike allocation strategies in a competitive dockless bike sharing market," *Journal of Cleaner Production*, vol. 233, pp. 869–879, 2019.

[45] P. Toth and D. Vigo, *The vehicle routing problem*. SIAM, 2002.

[46] G. B. Dantzig and J. H. Ramser, "The truck dispatching problem," *Management science*, vol. 6, no. 1, pp. 80–91, 1959.

[47] G. Laporte, "Fifty years of vehicle routing," *Transportation science*, vol. 43, no. 4, pp. 408–416, 2009.

[48] M. M. Solomon and J. Desrosiers, "Survey paper—time window constrained routing and scheduling problems," *Transportation science*, vol. 22, no. 1, pp. 1–13, 1988.

[49] J. Desrosiers, Y. Dumas, M. M. Solomon, and F. Soumis, "Time constrained routing and scheduling," *Handbooks in operations research and management science*, vol. 8, pp. 35–139, 1995.

[50] M. W. Savelsbergh and M. Sol, "The general pickup and delivery problem," *Transportation science*, vol. 29, no. 1, pp. 17–29, 1995.

[51] S. N. Parragh, K. F. Doerner, and R. F. Hartl, "A survey on pickup and delivery problems: Part ii: Transportation between pickup and delivery locations," *Journal für Betriebswirtschaft*, vol. 58, no. 2, pp. 81–117, 2008.

[52] B. L. Golden, L. Levy, and R. Vohra, "The orienteering problem," *Naval Research Logistics (NRL)*, vol. 34, no. 3, pp. 307–318, 1987.

[53] D. Feillet, P. Dejax, and M. Gendreau, "Traveling salesman problems with profits," *Transportation science*, vol. 39, no. 2, pp. 188–205, 2005.

[54] D. Pecin, A. Pessoa, M. Poggi, and E. Uchoa, "Improved branch-cut-and-price for capacitated vehicle routing," *Mathematical Programming Computation*, vol. 9, no. 1, pp. 61–100, 2017.

[55] C. Archetti and M. G. Speranza, "A survey on matheuristics for routing problems," *EURO Journal on Computational Optimization*, vol. 2, no. 4, pp. 223–246, 2014.

[56] L. Caccetta and S. P. Hill, "An application of branch and cut to open pit mine scheduling," *Journal of global optimization*, vol. 27, no. 2, pp. 349–365, 2003.

[57] B. Crevier, J.-F. Cordeau, and G. Laporte, "The multi-depot vehicle routing problem with inter-depot routes," *European journal of operational research*, vol. 176, no. 2, pp. 756–773, 2007.

[58] A. Vallera, P. Nunes, and M. Brito, "Why we need battery swapping technology," *Energy Policy*, vol. 157, p. 112481, 2021.

[59] A. Che, W. Wang, X. Mu, Y. Zhang, and J. Feng, "Tabu-based adaptive large neighborhood search for multi-depot petrol station replenishment with open inter-depot routes," *IEEE Transactions on Intelligent Transportation Systems*, vol. 24, no. 1, pp. 316–330, 2022.

[60] B. Pan, Z. Zhang, and A. Lim, "Multi-trip time-dependent vehicle routing problem with time windows," *European Journal of Operational Research*, vol. 291, no. 1, pp. 218–231, 2021.

[61] W. Wang and J. Zhao, "Partial linear recharging strategy for the electric fleet size and mix vehicle routing problem with time windows and recharging stations," *European Journal of Operational Research*, vol. 308, no. 2, pp. 929–948, 2023.

[62] Y. Wang, J. Zhou, Y. Sun, J. Fan, Z. Wang, and H. Wang, "Collaborative multidepot electric vehicle routing problem with time windows and shared charging stations," *Expert Systems with Applications*, vol. 219, p. 119654, 2023.

[63] J. Xiao, T. Zhang, J. Du, and X. Zhang, "An evolutionary multiobjective route grouping-based heuristic algorithm for large-scale capacitated vehicle routing problems," *IEEE transactions on cybernetics*, vol. 51, no. 8, pp. 4173–4186, 2019.

[64] Y. Cai, G. P. Ong, and Q. Meng, "Bicycle sharing station planning: From free-floating to geo-fencing," *Transportation research part C: emerging technologies*, vol. 147, p. 103990, 2023.

[65] M. Qi, W.-H. Lin, N. Li, and L. Miao, "A spatiotemporal partitioning approach for large-scale vehicle routing problems with time windows," *Transportation Research Part E: Logistics and Transportation Review*, vol. 48, no. 1, pp. 248–257, 2012.

[66] M. J. Santos, D. Jorge, T. Ramos, and A. Barbosa-Póvoa, "Green reverse logistics: Exploring the vehicle routing problem with deliveries and pickups," *Omega*, vol. 118, p. 102864, 2023.

[67] C. E. Miller, A. W. Tucker, and R. A. Zemlin, "Integer programming formulation of traveling salesman problems," *Journal of the ACM (JACM)*, vol. 7, no. 4, pp. 326–329, 1960.

[68] B. E. Gillett and L. R. Miller, "A heuristic algorithm for the vehicle-dispatch problem," *Operations research*, vol. 22, no. 2, pp. 340–349, 1974.

[69] M. L. Fisher, "Optimal solution of vehicle routing problems using minimum k-trees," *Operations research*, vol. 42, no. 4, pp. 626–642, 1994.

[70] O. Bräysy and M. Gendreau, "Vehicle routing problem with time windows, part i: Route construction and local search algorithms," *Transportation science*, vol. 39, no. 1, pp. 104–118, 2005.