

Bridging Policy and Real-World Dynamics: LLM-Augmented Rebalancing for Shared Micromobility Systems

Heng Tan¹, Hua Yan¹, Yu Yang¹

Abstract—Shared micromobility services such as e-scooters and bikes have become an integral part of urban transportation, yet their efficiency critically depends on effective vehicle rebalancing. Existing methods either optimize for average demand patterns or employ robust optimization and reinforcement learning to handle predefined uncertainties. However, these approaches overlook emergent events (e.g., demand surges, vehicle outages, regulatory interventions) or sacrifice performance in normal conditions. We introduce AMPLIFY, an LLM-augmented policy adaptation framework for shared micromobility rebalancing. The framework combines a baseline rebalancing module with an LLM-based adaptation module that adjusts strategies in real time under emergent scenarios. The adaptation module ingests system context, demand predictions, and baseline strategies, and refines adjustments through self-reflection. Evaluations on real-world e-scooter data from Chicago show that our approach improves demand satisfaction and system revenue compared to baseline policies, highlighting the potential of LLM-driven adaptation as a flexible solution for managing uncertainty in micromobility systems.

Index Terms—Shared Micromobility, Large Language Models, Planning, Intelligent Transportation Systems.

I. INTRODUCTION

Shared micromobility services (e.g., scooters and bikes) have rapidly expanded in urban areas around the world, offering flexible, low-emission, and last-mile transportation options for millions of users [1], [2]. According to the recent reports [3], global shared micromobility trips have surpassed billions annually, and cities are increasingly integrating such services into their public transit ecosystems. However, a key operational challenge in shared micromobility systems is vehicle rebalancing, i.e., ensuring that vehicles are distributed across space to meet spatiotemporal demand. Without effective rebalancing, certain regions may experience shortages while others face oversupply, resulting in poor user experience and revenue loss for operators [3]. In this work, we utilize electric scooters (e-scooters) as an example to study the rebalancing problem of shared micromobility vehicles.

In recent years, a variety of rebalancing strategies have been developed for shared micromobility systems. These methods can be broadly categorized into uncertainty-ignorant vehicle rebalancing [4], [4]–[7] and uncertainty-aware vehicle rebalancing [8]–[14]. Uncertainty-ignorant methods typically determine rebalancing actions based on the current

vehicle distribution and predictions of future demand prior to the next rebalancing interval. To optimize these decisions, researchers have employed techniques such as mixed-integer programming [15], [16] and sequential decision-making frameworks such as reinforcement learning [4], [6], [17], with the goal of maximizing system-level performance metrics (e.g., cumulative trip revenue or satisfied user demand). However, these approaches primarily emphasize optimality under general conditions or average demand patterns, while emergent or rare events are often overlooked, such as sudden demand surges [7], unexpected vehicle unavailability [3], or newly imposed city-level objectives [18]. To better address such uncertainties, prior studies have designed uncertainty-aware rebalancing methods based on robust optimization [8]–[10] and learning-based strategies [11], [12]. While these methods are effective in handling predefined uncertainties, they face an inherent trade-off between normal and emergent cases: the stronger the robustness to uncertainty, the more conservative the resulting policy becomes [8], [11], [12], [19], often leading to degraded performance under typical demand conditions. A practical solution is to introduce human intervention, as explored in related domains such as robotics [20], [21] and autonomous driving [22], [23], where humans can dynamically adapt policy outputs in emergent situations. However, deployment of such human-in-the-loop approaches in shared micromobility systems remains rare, due to limited human resources, scalability concerns, and the difficulty of integrating human inputs into automated decision-making pipelines.

Recent advances in large language model (LLM)-based agent frameworks have demonstrated their remarkable capabilities in situational reasoning and planning across various domains such as robotics [24], gaming [25], and autonomous systems [26]. For instance, LLM agents can be fine-tuned with experience datasets to adapt their actions dynamically to specific task objectives [27]. They can also leverage in-context learning and memory-driven reasoning to achieve generalization and scalability without fine-tuning [28]. Inspired by those works, we explore how an LLM-augmented agent can serve as a decision-making assistant to adapt vehicle rebalancing strategies towards the emergent situations.

We introduce AMPLIFY, an LLM-augmented policy adaptation framework for shared micromobility vehicle rebalancing. The framework is modular and consists of two decoupled components: (1) a vehicle rebalancing module that generates baseline strategies using either learning-based approaches (e.g., reinforcement learning) or rule-based methods (e.g., optimization); and (2) an LLM-augmented adaptation mod-

¹Heng Tan, Hua Yan, and Yu Yang are with the Department of Computer Science & Engineering, Lehigh University, Bethlehem, USA het221@lehigh.edu, huy222@lehigh.edu, yuyang@lehigh.edu

This work was supported in part by NSF grants 2246080, 2318697, and 2427915. We also thank the anonymous reviewers for their valuable comments and feedback.

ule that modifies these strategies in response to emergent conditions (e.g., sudden demand surges, event-driven disruptions, or new policy interventions from regulators). The adaptation module takes as input structured information, including the existing rebalancing strategy (i.e., baseline policy), current vehicle distribution, predicted demand, and a description of emergent conditions. It then produces an adapted rebalancing plan. To improve both reasoning quality and feasibility, we employ a self-reflection mechanism in which the LLM iteratively verifies, revises, and finalizes its initial adaptation. By decoupling the baseline policy from the adaptation module, our framework achieves high flexibility and compatibility with diverse rebalancing algorithms. It enables real-time, human-like intervention without retraining or encoding scenario-specific rules, offering a generalizable and scalable solution for adapting to uncertainty in shared micromobility systems.

In summary, the key contributions are as follows:

- We design a simple yet effective LLM-augmented policy adaptation framework for shared micromobility vehicle rebalancing under emergent scenarios. Unlike traditional approaches that rely on retraining or robust pre-planning, our framework enables real-time strategy adaptation with minimal computational overhead.
- We develop a modular system that decouples the vehicle rebalancing backbone (e.g., optimization- or learning-based) from the LLM-augmented agent. The agent reasons over system context and emergent conditions, generates adapted rebalancing strategies, and iteratively refines them through self-reflection.
- We conduct comprehensive evaluations using real-world e-scooter data from Chicago. Experimental results show that our framework substantially improves pre-optimized rebalancing strategies in emergent scenarios (e.g., sudden demand surges or vehicle outages), leading to higher system-level performance (e.g., demand satisfaction and revenue). Moreover, we find that grounding the LLM in baseline policies significantly enhances its ability to produce effective and stable adaptations, compared to generating strategies by LLMs directly.

II. RELATED WORK

A. Shared Micromobility Vehicle Rebalancing

A wide range of methods has been developed to optimize vehicle rebalancing strategies in shared micromobility systems, with the goal of maximizing system-level performance such as user satisfaction or operational efficiency. These methods can be broadly categorized into two groups: (i) some researchers assume stable and predictable operational environments and utilize non-uncertainty-aware approaches, including optimization-based methods [16], [29], [30] and RL-based methods [4], [6], [17], [31], [32]. For example, [4] proposes a regulatory multi-operator vehicle scheduling framework for shared electric scooters, which incorporates an interaction between a city regulator and multiple shared micromobility operators to balance the tradeoff between city-level goals and operators' interests. However, those methods,

which are trained on stable and predictable operational environments, struggle to adapt to these unforeseen scenarios, resulting in suboptimal performance during critical moments. (ii) Other researchers employ uncertainty-aware approaches, which explicitly account for disruptions and variations in demand or supply [8], [9], [12], [33]. For example, [8] develops a distributionally robust optimization approach for vehicle rebalancing in AMoD systems under deep demand uncertainty, explicitly modeling worst-case scenarios to improve resilience. However, those methods generally lack the flexibility to adapt to real-time emergent situations, where unexpected demand surges, vehicle failures, or new city-level constraints require immediate and context-aware adaptations to the existing scheduling policy.

B. LLM Agents for Planning

Recent advances in large language model (LLM)-based agent frameworks have demonstrated their remarkable capabilities in situational reasoning and adaptive planning in complex and dynamic environments [24]–[27]. For example, [24] provides a modular planning framework that enhances LLM tool execution accuracy by structuring multi-step workflows. [26] introduces closed-loop feedback for dynamic plan refinement, enabling more robust adaptation over long horizons. However, applying LLMs to transportation systems or vehicle scheduling under emergent situations remains largely underexplored. Our work fills this gap by leveraging an LLM-based agent to adapt pre-optimized vehicle rebalancing strategies in response to emergent situations, demonstrating its potential in real-time operational adaptation.

III. DESIGN

A. Problem Description

Problem Setting: We consider an urban city divided into N regions based on official community boundaries [18], and partition the day into T equal-length time intervals. Let S_t^i denote the number of available shared micromobility vehicles in region i at the beginning of time slot t for $1 \leq i \leq N$. User demand is defined over origin-destination (OD) pairs. Specifically, $U_t^{i,j}$ denotes the number of trip requests from region i to region j during time slot t for $1 \leq i, j \leq N$. Therefore, the overall vehicle distribution and user demand across the city at the beginning of time slot t are denoted respectively as $S_t \in \mathbb{N}^N$, and $U_t \in \mathbb{N}^{N \times N}$.

Rebalancing: To meet future user demand, it is essential for the system operator to redistribute or rebalance shared micromobility vehicles to different areas. We define $a_t \in \mathbb{N}^{N \times N}$ ($a_t = \{a_t^{i,j}\}_{1 \leq i, j \leq N}$) as the operator's rebalancing strategy at the beginning of the timeslot t . In our work, we modularize the vehicle rebalancing generation, which inputs the current vehicle distribution and predicted user demand of future h time slots. Therefore, the rebalancing strategy a_t is formulated as:

$$a_t \leftarrow f_{reb}(S_t, U_{t:t+h}), \quad (1)$$

where f_{reb} is the vehicle rebalancing policy, which is assumed to be pre-determined by existing works [17], [34].

Adaptation under Emergent Scenarios: Although the rebalancing strategy a_t is generated by the pre-determined vehicle rebalancing policy f_{reb} , real-world operations often encounter emergent scenarios that violate prior assumptions, such as sudden demand surges, vehicle malfunctions, or regulatory interventions (e.g., constraints). We define each emergent scenario at time t as a tuple $\mathcal{E}_t = \{\Delta U_t, \Delta S_t, C_t\}$, which respectively describes unexpected variations in user demand, vehicle availability, and constraint conditions imposed by the external environment. To enhance system robustness under such emergent scenarios, an adaptation is needed to update the rebalancing strategy based on real-time observations. Specifically, we define an adaptation policy π_{adj} that takes the initial rebalancing strategy a_t , the current vehicle state S_t , the predicted demand of future h time slots $U_{t:t+h}$, and the emergent situation \mathcal{E}_t as input, and outputs a adapted vehicle rebalancing strategy \hat{a}_t :

$$\hat{a}_t \leftarrow \pi_{adj}(a_t, S_t, \mathcal{E}_t). \quad (2)$$

Objective: The goal of adaptation is to ensure that vehicle rebalancing strategies remain effective in the presence of sudden distribution shifts (e.g., demand surges or vehicle unavailability) and evolving system objectives (e.g., a dynamic system-level goal such as maximizing vehicle deployment equity). Therefore, we formalize the adaptation objective as maximizing the overall system utility under emergent situations:

$$\max_{\pi_{adj}} \mathbb{E}_{\mathcal{E}_t \sim \mathcal{P}_{\mathcal{E}}} [\mathcal{R}(\hat{a}_t, \mathcal{E}_t)], \quad (3)$$

where $\mathcal{P}_{\mathcal{E}}$ is the probability distribution of emergent situations. \mathcal{R} quantifies system-level metrics such as demand satisfaction, system revenue, and vehicle deployment equity under adapted strategies.

B. Design Overview

We design a modular framework for adaptive shared micromobility vehicle rebalancing, which integrates pre-determined vehicle rebalancing strategies with an LLM-augmented agent. As illustrated in Fig. 1, the system consists of three key components: (1) **Environment:** It simulates the shared micromobility system operations, including vehicle dynamics, user demand, and request fulfillment. (2) **Vehicle Rebalancing Module:** It generates initial rebalancing strategies based on current vehicle distribution and predicted future user demand across all regions. This module can adopt either optimization-based or learning-based algorithms and operates independently of downstream adaptation logic. (3) **LLM-augmented Adaptation Agent:** When an emergent situation occurs, it is described in natural language and passed to the LLM agent, together with the current strategy, system-level status (i.e., vehicle distribution and predicted future user demand), and goals (e.g., system revenue). The LLM processes this information and outputs an adapted vehicle rebalancing strategy for execution. To enable more stable and feasible adaptations, there is a self-reflection mechanism that prompts the LLM-augmented agent to verify, revise, and finalize its initial adaptation iteratively.

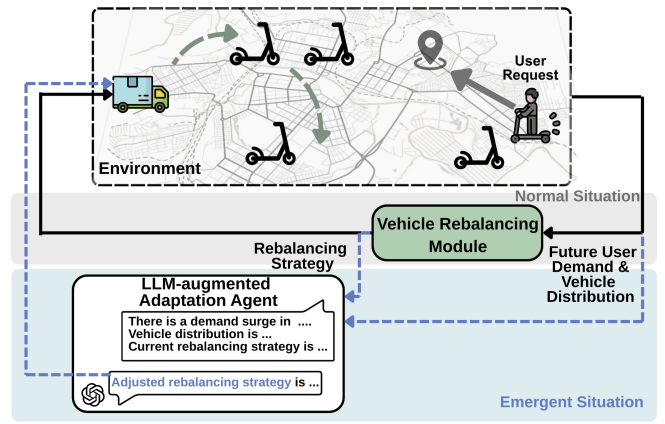


Fig. 1: LLM-augmented policy adaptation framework for shared micromobility vehicle rebalancing under emergent scenarios

C. Natural Language Description of Emergent Situations

In real-world shared micromobility systems, emergent situations—such as sudden demand surges, vehicle unavailability, or dynamic goals—are typically communicated to operators via unstructured textual reports [35], [36]. These include messages like “a large crowd is forming near the stadium”, “vehicles in Region 2 are out of service due to maintenance”, or “ensure better service equity across southern districts”. Unlike structured telemetry, these natural language inputs capture real-time, contextualized, and often ambiguous human insights. To reflect this reality, we adopt *natural language* as the input modality for emergent situations. This enables our LLM-augmented agent to leverage its strong capabilities in understanding and reasoning over natural language inputs, allowing for flexible interpretation and transformation into actionable internal representations. We categorize emergent situations into three major types: (1) **Rising Demand: Sudden Surges in User Requests.** Demand surges represent one of the most frequent and critical emergent situations in shared micromobility systems [3], [37]. In real-world scenarios, however, such surges are not always directly measurable. To comprehensively assess the robustness of our LLM-augmented agent, we consider two subtypes: (i) *observable surges*, where the magnitude of the increased demand are explicitly available (e.g., from sensors or early warnings), and (ii) *latent surges*, where only vague signals are conveyed via natural language (e.g., “a large crowd is forming near the stadium”) without precise quantification. These two cases are used to evaluate how well the LLM-augmented agent can infer and adapt to demand shifts under varying degrees of information completeness. (2) **Shrinking Supply: Vehicle Shortages and Failures.** Vehicle unavailability frequently arises due to maintenance, vandalism, or battery depletion, especially in shared micromobility systems with large, distributed fleets [3], [38]. These disruptions are often reported by region staff or monitoring systems in natural language, such as “bikes in Region 3 are under maintenance.” Therefore, we introduce vehicle unavailability scenarios to test whether the LLM agent can recognize unavailable supply, avoid infeasible reallocations,

and adapt rebalancing accordingly.

(3) **Dynamic Goals.** In addition to maximizing user demand satisfaction, real-world shared micromobility systems are increasingly expected to fulfill additional system-level goals imposed by regulators [4], [18], such as equity, coverage, or sustainability. In this work, we use vehicle deployment equity as an example, which is ensuring a fair spatial distribution of the demand-supply ratio across regions. Notably, this natural language framing enables the flexible definition and adaptation of equity goals under various metrics (e.g., entropy [39], variance [4], minimum coverage [18]), and can be extended to other system-level goals in future applications.

D. LLM-augmented Rebalancing Strategy Adaptation

To enable dynamic adaptation of rebalancing strategies in response to emergent scenarios, we employ an LLM as a zero-shot planner that refines pre-generated vehicle rebalancing decisions. Compared with directly using LLMs for planning, our experiments show that incorporating pre-generated strategies significantly improves planning effectiveness and stability. The LLM is provided with a structured prompt that includes a comprehensive system background, real-time system status, initial vehicle rebalancing strategy, and explicit instructions as shown in Box 1. The goal is to generate an adapted rebalancing plan that is both operationally feasible and aligned with emergent requirements. The prompt is composed of several key sections: (1) **Background:** To guide the LLM’s adaptation reasoning, we formulate a structured background prompt that specifies the system state (vehicle availability V , demand prediction D), the operator’s current rebalancing plan R , and the global objective of maximizing net revenue. This structured grounding provides the LLM with the necessary operational context to evaluate trade-offs and generate effective rebalancing adaptations. (2) **Real-time Context Input:** To enable the LLM to make informed and context-aware adaptations, we present a natural language encoding of the current system status and the operator’s initial vehicle rebalancing plan. Specifically, we include: (i) region-level vehicle supply, (ii) predicted demand flows across origin-destination pairs, and (iii) region-level trip variance statistics derived from historical data. The initial rebalancing strategy generated by the pre-determined vehicle rebalancing policy is also described in natural language form. (3) **Adaptation Instruction:** To guide the LLM in making principled and goal-aware rebalancing decisions, we formulate a multi-step adaptation instruction. This instruction prompts the LLM to think through the structured reasoning steps before proposing any adaptation to the initial rebalancing plan, including region-level rebalancing necessity, goal alignment analysis, and adaptation justification.

To further enhance the robustness, feasibility, and multi-goal alignment of the LLM-generated rebalancing strategy adaptation, we incorporate a self-reflection loop that enables the LLM to critically evaluate and revise its own output, inspired by recent advances in reflective prompting and chain-of-thought self-verification [26], [28]. The self-reflection loop proceeds as follows: (1) The LLM receives

a full prompt containing system background, current vehicle distribution, initial rebalancing plan, predicted user demand, emergent directive, and a prior adaptation output. (2) It is instructed to reassess this output by examining key feasibility and goal-alignment criteria, including dimensional validity, conservation constraint, and task satisfaction.

Box 1: Prompt for LLM-augmented Rebalancing

Background: You are an intelligent agent embedded in a shared micromobility system. The city is divided into N regions... // problem formulation of shared micromobility scheduling

System Status:

- **System Time:** 00:00
- **Current Vehicle Distribution** ($S_t \in \mathbb{N}^N$): [12, 5, 8, 3, ..., 7] // Each entry S_t^i denotes the number of vehicles in region i
- **Predicted future Demand** ($U_{t:t+h} \in \mathbb{N}^{N \times N}$): Region 1 \rightarrow 2: 10, Region 3 \rightarrow 5: 6, ... // Each entry $U_t^{i,j}$ denotes the number of predicted trips from region i to region j
- **Region-level trip variance:** For each region, the trip variance is provided based on historical trip data: region 1 = {avg: 8.2, std: 3.5, min: 2, max: 13}, ... // statistical patterns of trip demand in each region

Initial Rebalancing Strategy (a_t): Move 2 vehicles from Region 4 to 1, 3 from Region 6 to 3, ... // Generated by a pre-determined rebalancing policy

Emergent Situation: The city regulator requests the vehicle deployment equity across all regions...

Instruction: Your goal is to adapt the operator’s vehicle rebalancing strategy under the current emergent situation, considering the provided real-time vehicle distribution, predicted future demand, and historical region-level trip variance. You must think through the following steps before proposing any adaptation:

- 1. Determine whether reallocation is needed and from which regions to meet both the requirements of the emergency and the operator.
- 2. Explain how you balance revenue maximization against the requirement of the emergency. Please calculate the goal achievements before and after your adaptation if the goal of the emergency is different from that of the operator.
- 3. Justify your adaptation with clear reasoning.

Constraint: The total vehicle number in the city after your adaptation must be consistent with that before your adaptation.

Output Example: An example of output format...

(3) The LLM then produces a refined adaptation strategy, which is injected back into the next round of reflection as the new candidate solution. This self-refinement process improves adaptation quality without requiring external super-

vision. Empirically, we find that this iterative self-reflection often corrects common issues such as over-allocation, infeasible transfers, or failure to balance multiple goals. The final output is typically more robust, constraint-consistent, and aligned with evolving system requirements.

IV. EVALUATION

A. Experiment Setting

Implementation: We conduct our experiments on a publicly available real-world shared e-scooter dataset from Chicago [18], comprising over 629,000 e-scooter trips operated by Lime, Spin, and Bird between June and September 2022. The dataset contains information such as trip time, distance, operator ID, departure time, and region, among other attributes. The entire city is partitioned into 77 regions based on the existing community divisions in the dataset. A day is divided into 24 hours, with operators’ rebalancing occurring every 12 hours. For users’ vehicle selections, only nearby vehicles with adequate energy can meet user demand.

We implement our method and baselines with PyTorch 1.9.1, Python-mip 1.14.2, and gym 0.21.0 in a Python 3.7 environment, and we train them on a server equipped with 32 GB of memory and a GeForce RTX 3080 Ti GPU. For LLMs, we choose GPT-4o as our LLM-augmented agent. The maximum number of iterations in self-reflection is set as 10. Hyperparameters for vehicle rebalancing baselines are fine-tuned based on the range in the original papers.

Baselines: We evaluate the performance of our model AMPLIFY with the following vehicle rebalancing baselines: (1) **SDSM**. It is a rule-based demand-supply matching method where each operator rebalances the vehicles to each region based on the ratio of the historical demand. (2) **GA** [34]. It is an optimization-based method that utilizes a genetic algorithm to find the optimal vehicle scheduling strategies to maximize the demand satisfaction rate. (3) **RECOMMEND** [17]. It is a state-of-the-art shared electric micromobility vehicle scheduling algorithm that utilizes a MARL-based method to learn the optimal vehicle scheduling policies to maximize the system service performance. (4) **LLM-planning**. It is an LLM-augmented agent that directly makes vehicle rebalancing decisions. We adopt the Box 1 prompt without initial rebalancing strategies.

Metrics: The evaluation metrics are as follows: (1) **Average satisfaction rate:** The average satisfaction rate represents the average ratio of satisfied demand to the total user demand among all the regions. (2) **Vehicle usage equity** [4]: the metric measures the difference between demand-supply ratios of individual regions and the whole city.

B. Overall Performance

To comprehensively evaluate the adaptability of our LLM-augmented framework, we randomly select ten days in two months as emergent dates and design three distinct types of emergent scenarios that reflect real-world disruptions and policy shifts to evaluate the average performance of different methods under the following different emergent situations: (1) **Rising Demand:** A sudden surge in user demand

across the city, simulated by scaling region-level demand with varying rising ratios ranging from 20% to 100%. (2) **Shrinking Supply:** A reduction in the total available vehicle supply, representing supply-side constraints such as ongoing maintenance. We simulate this by decreasing the available fleet by 5% to 20% (randomly select vehicles across all regions). (3) **Dynamic Goals:** In this setting, the regulator imposes an additional constraint to improve equity in vehicle deployment. We compare our method with other baselines in these three scenarios, and the experiment results (containing average performance with its variance) are shown in Table I. In our work, our model is the LLM-augmented adaptation of the RECOMMEND rebalancing strategy, and we discuss the effectiveness of underlying rebalancing strategies in Section IV-C.1.

Based on this table, we can know that **(1) Rising Demand:** as user demand increases from 20% to 100%, all baseline strategies exhibit a steady decline in satisfaction rate due to increasing service pressure. Notably, AMPLIFY maintains a remarkably stable performance (from 97.51% to 97.06%), significantly outperforming all baselines. For example, at 100% increased demand, AMPLIFY achieves a satisfaction rate of 97.06%, compared to 90.62% for LLM-planning and only 63.12% for SDSM. This result illustrates the effectiveness of our LLM-augmented agent in leveraging initial strategies and promptly redistributing vehicles in response to sudden demand surges. **(2) Shrinking Supply:** Even though AMPLIFY still outperforms all baselines in this scenario, its performance stability reduces compared with that in the scenario of rising demand. The possible reason is the random removal of vehicles across the city, which breaks the original spatial supply distribution and increases the complexity of accurate adaptation for the initial vehicle rebalancing strategy. In contrast, the rising demand scenario benefits from more predictable spatiotemporal patterns, allowing the LLM to generate more effective adaptations even under scaling stress. **(3) Dynamic Goals:** In the equity-focused setting, both the LLM-augmented agent and the LLM-planning baseline (i.e., directly generating rebalancing strategies without any underlying strategy) outperform traditional baselines. This demonstrates that LLMs can not only understand natural language descriptions of complex policy objectives (e.g., minimizing service inequity) but also incorporate these goals into the adaptation of vehicle rebalancing strategies by step-by-step reasoning.

C. Ablation Study

1) *The effectiveness of underlying rebalancing strategies:* To investigate how the choice of underlying rebalancing strategy affects the performance of the LLM-augmented agent, we conduct an ablation study across three representative scenarios: (i) rising demand (+50%), (ii) shrinking supply (-10%), and (iii) a dynamic goal (i.e., maximizing vehicle deployment equity). In each scenario, we compare our method with the other three baselines: (i) LLM-planning, (ii) LLM-GA: the LLM adapts vehicle rebalancing strategies generated by a classical genetic algorithm [34], and (iii)

TABLE I: Performance comparison (i.e., demand satisfaction rate (%) and equity) of different vehicle rebalancing methods under different emergent scenarios. AMPLIFY denotes our LLM-augmented adaptation of the RECOMMEND rebalancing strategy.

Method	Rising Demand (\uparrow)				Shrinking Supply (\uparrow)				Equity (\uparrow)
	20%	50%	80%	100%	5%	10%	15%	20%	
SDSM	84.12 (± 0.44)	73.36 (± 0.45)	66.61 (± 0.47)	61.45 (± 0.49)	81.10 (± 0.44)	74.07 (± 0.47)	66.68 (± 0.49)	57.01 (± 0.53)	-164.27 (± 2.41)
GA	87.94 (± 0.79)	81.76 (± 0.81)	75.29 (± 0.82)	70.46 (± 0.83)	85.97 (± 0.77)	83.76 (± 0.80)	80.86 (± 0.84)	76.35 (± 0.87)	-111.54 (± 2.59)
RECOMMEND	93.82 (± 1.34)	88.47 (± 1.35)	81.57 (± 1.37)	80.68 (± 1.41)	91.61 (± 1.33)	89.04 (± 1.34)	84.92 (± 1.38)	80.04 (± 1.43)	-81.34 (± 2.92)
LLM-planning	91.69 (± 2.95)	91.50 (± 2.96)	91.07 (± 2.97)	90.86 (± 2.99)	84.56 (± 2.04)	82.65 (± 2.06)	81.38 (± 2.11)	79.85 (± 2.16)	-76.38 (± 5.74)
AMPLIFY (ours)	97.34 (± 1.46)	97.22 (± 1.47)	97.10 (± 1.47)	96.87 (± 1.49)	94.93 (± 1.35)	93.79 (± 1.38)	91.40 (± 1.41)	89.38 (± 1.46)	-51.79 (± 1.84)
Average Improvement	8.89%	16.05%	23.48%	27.69%	10.63%	13.85%	16.49%	21.92%	52.22%







Fig. 2: Performance under rising demand

Fig. 3: Performance under shrinking supply

Fig. 4: Performance under dynamic goals

Fig. 5: Performance under different scenario descriptions

LLM-SDSM: the LLM adapts vehicle rebalancing strategies generated by a rule-based method (i.e., SDSM). Each method in each scenario is evaluated by prompting the LLM multiple times to capture the variability in its outputs. We visualize the performance variance of different methods using boxplots, as shown in Fig. 2, Fig. 3, and Fig. 4.

We observe that the effectiveness of the LLM-augmented adaptation agent is highly dependent on the quality of the initial vehicle rebalancing strategy. When provided with a strong baseline (e.g., the RL-based method RECOMMEND), the LLM-augmented adaptation performs exceptionally well, adapting the strategy with high precision and stability. In contrast, when the initial strategy is weak (e.g., the rule-based method SDSM), the LLM-augmented agent struggles to overcome fundamental planning deficiencies, leaving substantial room for improvement. Moreover, in the LLM-planning setting, where no initial strategy is provided, we observe substantial performance variance across all scenarios compared with other baselines that are provided with initial vehicle rebalancing strategies. It highlights the difficulty of open-ended planning tasks for LLMs, especially in constrained domains like vehicle allocation, and indirectly demonstrates that the quality of the initial strategy heavily influences the effectiveness of LLM adaptation.

2) *The effectiveness of different LLMs*: To evaluate how different LLMs perform in adapting pre-determined vehicle rebalancing strategies to various emergent scenarios, we conduct an ablation study comparing four adaptation models using different language backbones: GPT-3.5-turbo, LLaMA-3.1-8B, LLaMA-3.1-70B, and GPT-4o (used in AMPLIFY) to adapt vehicle rebalancing strategies generated by RECOMMEND, as shown in Table II.

Based on the experiment results, we can know that (i) The strategy adaptation capability is tied to the model’s reasoning and comprehension capacity: Smaller models such as GPT-3.5 and LLaMA-3.1-8B often fail to generate ef-

fective strategy adaptations, generating strategies worse than the original baseline RECOMMEND. In contrast, larger models like LLaMA-3.1-70B and GPT-4o successfully adapt the initial strategy, consistently yielding improved outcomes across all scenarios. (ii) While both larger models (i.e., LLaMA-3.1-70B and GPT-4o) exhibit strong adjustment capability, GPT-4o (AMPLIFY) delivers the highest improvements, particularly in the scenario of dynamic goals (i.e., equity). This suggests that GPT-4o possesses stronger reasoning and generalization ability when integrating complex system-level goals into its adaptations.

TABLE II: The performance comparison of the adaptation model with different LLMs

Method	Rising Demand	Shrinking Supply	Equity
GPT-3.5-turbo	85.51%	85.56%	-104.97
Llama-3.1-8B	85.35%	88.74%	-98.12
Llama-3.1-70B	93.27%	90.17%	-78.14
GPT-4o	97.37%	93.28%	-50.48

3) *The effectiveness of self-reflection*: We find that the self-reflection mechanism is critical for ensuring both the correctness and effectiveness of the LLM-augmented rebalancing strategies. In an ablation study where the self-reflection loop is disabled, the LLM frequently outputs invalid or flawed strategies even when the prompt explicitly encodes key operational constraints. These include: (i) malformed output formats that cannot be parsed into action vectors, (ii) mismatches in total vehicle count before and after adaptation, (iii) violations of basic constraints (e.g., negative vehicle allocations), and (iv) disregard for critical input signals such as emergency needs or equity goals. To quantify the effectiveness of the self-reflection module, we conduct 20 independent runs of the adaptation process under both settings. Without self-reflection, only 3 out of 20 runs (15%) yielded valid and usable rebalancing strategies. In contrast, with self-reflection enabled, 17 out of 20 runs

(85%) successfully produced valid, constraint-satisfying, and context-aware strategies. The experiment results demonstrate that self-reflection acts as a crucial internal verification step, guiding the LLM to revise and validate its reasoning before finalizing an adapted vehicle rebalancing decision. Nevertheless, we acknowledge that in a small portion (15%) of cases, the self-reflection mechanism still fails to converge to a valid adaptation. To address these cases and ensure robustness in real deployments, we provide two backup plans: (i) Increase the iteration depth of the self-reflection loop; (ii) Fallback to the initial strategy if no valid adaptation is produced.

D. Impact of Factor

1) *The impact of explicit information in emergent scenarios:* To examine the importance of informing the LLM about the specific surge magnitude in emergent demand scenarios, we conduct an ablation study by comparing two variants of our adaptation agent: one that is explicitly prompted with the region-wise rising demand ratio (Known), and one that is only informed of the region-wise rising signal without any surge magnitude information (Unknown).

Based on the experimental results in Fig. 5, we observe that: (i) even without explicit surge magnitude, the LLM can produce meaningful adaptations by leveraging its internal assumptions. This is because the prompt includes region-level demand variability statistics (mean, standard deviation, min, and max), which implicitly guide the LLM to hypothesize a reasonable demand shift. This capability is evident at the 20% rising point, where the satisfaction rates for known and unknown variants are nearly identical. (ii) Across all rising levels, the known variant consistently outperforms the unknown variant. Not only does it achieve higher satisfaction rates, but it also maintains a more stable performance curve with a lower degradation slope as the rising ratio increases. This indicates that providing precise, structured surge information enables the LLM to generate more targeted and effective rebalancing strategies. (iii) The performance degradation of the Known variant reflects the structural limitation of finite vehicle supply: regardless of the LLM’s awareness of surge magnitude, there is a natural bottleneck to how much adapted vehicle rebalancing strategies can serve extreme demand increases.

2) *The impact of different equity definitions:* To investigate the flexibility of our framework under different equity definitions, we compare two widely used fairness metrics: the Gini coefficient [40] and the Theil index [39]. Both metrics quantify the inequality in the spatial distribution of service. They differ in sensitivity and formulation. While the Gini coefficient emphasizes rank-based inequality, the Theil index captures entropy-based disparity and is more sensitive to deviations in the tails.

As shown in Table III, prompting the LLM with different equity definitions enables our adaptation agent to generate rebalancing strategies that consistently reduce regional inequality compared with the initial strategies. This demonstrates that the LLM can interpret and adapt to multiple equity formulations expressed in natural language, producing more

equitable rebalancing outcomes without retraining. These results highlight the flexibility and effectiveness of our model in aligning outputs with diverse equity objectives.

TABLE III: The impact of equity definition

Method	Gini coefficient (\downarrow)	Theil index (\downarrow)
RECOMMEND	0.6738 (± 0.001)	3.9702 (± 0.001)
AMPLIFY	0.5007 (± 0.022)	3.9213 (± 0.002)

E. Case Study

To evaluate our LLM-augmented vehicle rebalancing strategy under real-world emergent conditions, we select September 5th, 2022 (Monday) in Chicago as a test day. This date coincides with multiple large events, such as Labor Day, the Bitter Jester Music Festival, African Festival of the Arts, and the Taste of Polonia Festival, all contributing to atypical user mobility patterns. On this day, the system recorded 12,373 shared micromobility trips, which is nearly double the average Monday volume of 6,603 trips. Fig. 6 presents the demand distribution across all regions for the Monday average and the selected study day, highlighting regions with unusually high demand surges (e.g., region 5 and region 10).

To test whether our model can effectively adapt vehicle rebalancing decisions to such a real-world emergent scenario, we assume that we have prior knowledge of these rising-demand regions before users initiate trips and prompt the context to the LLM-augmented agent to generate the adapted vehicle rebalancing strategy. This is a reasonable assumption since the locations of these events are known. We compare the performance of the adapted vehicle rebalancing strategy with the initial strategy provided by RECOMMEND. The experiment result in Fig. 7 shows that the baseline underperforms on this emergent day, achieving only 83.44% fulfillment, which is well below its monthly average of 90.61%. In contrast, our work achieves a significant performance improvement, with a rise to 92.87%, which demonstrates the potential of our LLM-augmented adaptation framework to adapt rebalancing strategies in real-world emergent settings.




Fig. 6: The comparison of two demand distributions




Fig. 7: The performance comparison of two methods

V. CONCLUSION

In this work, we focus on the problem of adapting shared micromobility vehicle rebalancing strategies under various emergent scenarios. We design an LLM-augmented adaptation framework that leverages context-aware prompts and self-reflection to refine vehicle rebalancing strategies provided by pre-determined policies. The evaluation results show that our method consistently improves performance across various settings, achieving at least an 8.89% improvement in satisfaction rate and 52.22% in vehicle deployment equity.

REFERENCES

- [1] S. Zhong, W. Yubeaton, W. Lyu, G. Wang, D. Zhang, and Y. Yang, "Rlife: Remaining lifespan prediction for e-scooters," in *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, pp. 3544–3553, 2023.
- [2] S. Zhong, H. Wei, W. Lyu, G. Yang, Z. Hong, G. Wang, Y. Yang, and D. Zhang, "Adatrans: Adaptive transfer time prediction for multimodal transportation modes," in *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, pp. 3443–3452, 2024.
- [3] National Association of City Transportation Officials (NACTO), "2023 shared micromobility in the U.S. and Canada report," July 2024.
- [4] H. Tan, H. Yan, Y. Yuan, G. Wang, and Y. Yang, "Realism: A regulatory framework for coordinated scheduling in multi-operator shared micromobility services," in *Proceedings of the 33rd ACM International Conference on Advances in Geographic Information Systems*, pp. 894–904, 2025.
- [5] Y. Wu, T. Liu, and B. Du, "Fleet sizing and static rebalancing strategies for shared e-scooters: A case study in indianapolis, usa," *Transportation Research Part A: Policy and Practice*, vol. 190, p. 104287, 2024.
- [6] S. He and K. G. Shin, "Socially-equitable interactive graph information fusion-based prediction for urban dockless e-scooter sharing," in *Proceedings of the ACM Web Conference 2022*, pp. 3269–3279, 2022.
- [7] H. Tan, H. Yan, L. Yang, and Y. Yang, "Small fleet, big impact: Enhancing shared micromobility efficiency through minimal autonomous vehicle deployment," in *Proceedings of the 12th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*, pp. 203–212, 2025.
- [8] X. Li, X. Guo, Q. Wang, G. Zardini, and J. Zhao, "Robust vehicle rebalancing with deep uncertainty in autonomous mobility-on-demand systems," *arXiv preprint arXiv:2507.04520*, 2025.
- [9] Y. Zhang, J. Cheng, O. Cats, and S. S. Azadeh, "Stochastic optimization under supply uncertainty for multimodal trip planning based on demand prediction," 2025.
- [10] B. P. Bruck, W. P. Coutinho, and P. Munari, "The robust bike sharing rebalancing problem: Formulations and a branch-and-cut algorithm," *European Journal of Operational Research*, vol. 325, no. 1, pp. 67–80, 2025.
- [11] S. He, S. Han, and F. Miao, "Robust electric vehicle balancing of autonomous mobility-on-demand system: A multi-agent reinforcement learning approach," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 5471–5478, IEEE, 2023.
- [12] S. He, Y. Wang, S. Han, S. Zou, and F. Miao, "A robust and constrained multi-agent reinforcement learning electric vehicle rebalancing method in amod systems," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 5637–5644, IEEE, 2023.
- [13] H. Yan, H. Tan, H. Wang, D. Zhang, and Y. Yang, "Robust route planning under uncertain pickup requests for last-mile delivery," in *Proceedings of the ACM Web Conference 2024*, pp. 3022–3030, 2024.
- [14] Q. Yang, Z. Xie, H. Wei, D. Zhang, and Y. Yang, "Mallight: Influence-aware coordinated traffic signal control for traffic signal malfunctions," in *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, pp. 2879–2889, 2024.
- [15] G. Guo and T. Xu, "Vehicle rebalancing with charging scheduling in one-way car-sharing systems," *IEEE Transactions on Intelligent Transportation Systems*, 2020.
- [16] Y. Yuan, D. Zhang, F. Miao, J. Chen, T. He, and S. Lin, "p²charging: proactive partial charging for electric taxi systems," in *2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS)*, pp. 688–699, IEEE, 2019.
- [17] H. Tan, Y. Yuan, S. Zhong, and Y. Yang, "Joint rebalancing and charging for shared electric micromobility vehicles with energy-informed demand," in *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, pp. 2392–2401, 2023.
- [18] Chicago, "E-scooter trips - 2020 - chicago data portal," 2020.
- [19] F. Miao, S. He, L. Pepin, S. Han, A. Hendawi, M. E. Khalefa, J. A. Stankovic, and G. Pappas, "Data-driven distributionally robust optimization for vehicle balancing of mobility-on-demand systems," *ACM Transactions on Cyber-Physical Systems*, vol. 5, no. 2, pp. 1–27, 2021.
- [20] H. Liu, S. Nasiriany, L. Zhang, Z. Bao, and Y. Zhu, "Robot learning on the job: Human-in-the-loop autonomy and learning during deployment," *The International Journal of Robotics Research*, p. 02783649241273901, 2022.
- [21] J. Luo, C. Xu, J. Wu, and S. Levine, "Precise and dexterous robotic manipulation via human-in-the-loop reinforcement learning," *Science Robotics*, vol. 10, no. 105, p. eads5033, 2025.
- [22] J. Wu, Z. Huang, C. Huang, Z. Hu, P. Hang, Y. Xing, and C. Lv, "Human-in-the-loop deep reinforcement learning with application to autonomous driving," *arXiv preprint arXiv:2104.07246*, 2021.
- [23] Z. Huang, Z. Sheng, C. Ma, and S. Chen, "Human as ai mentor: Enhanced human-in-the-loop reinforcement learning for safe and efficient autonomous driving," *Communications in Transportation Research*, vol. 4, p. 100127, 2024.
- [24] G. Zeng, X. Chen, J. Hu, S. Qi, Y. Mao, Z. Wang, Y. Nie, S. Li, Q. Feng, P. Qiu, et al., "Routine: A structural planning framework for llm agent system in enterprise," *arXiv preprint arXiv:2507.14447*, 2025.
- [25] H. Tan, H. Yan, and Y. Yang, "Llm-guided reinforcement learning: Addressing training bottlenecks through policy modulation," *arXiv preprint arXiv:2505.20671*, 2025.
- [26] H. Sun, Y. Zhuang, L. Kong, B. Dai, and C. Zhang, "Adaplaner: Adaptive planning from feedback with language models," *Advances in neural information processing systems*, vol. 36, pp. 58202–58245, 2023.
- [27] W. Yao, S. Heinecke, J. C. Niebles, Z. Liu, Y. Feng, L. Xue, R. Murthy, Z. Chen, J. Zhang, D. Arpit, et al., "Retroformer: Retrospective large language agents with policy gradient optimization," *arXiv preprint arXiv:2308.02151*, 2023.
- [28] N. Shinn, F. Cassano, A. Gopinath, K. Narasimhan, and S. Yao, "Reflexion: Language agents with verbal reinforcement learning," *Advances in Neural Information Processing Systems*, vol. 36, pp. 8634–8652, 2023.
- [29] G. Lee, J. S. Lee, and K. S. Park, "Battery swapping, vehicle rebalancing, and staff routing for electric scooter sharing systems," *Transportation Research Part E: Logistics and Transportation Review*, vol. 186, p. 103540, 2024.
- [30] J. Yun, J. Lee, and J. Kim, "Automated mobility-on-demand service improvement strategy through latent class analysis of stated preference survey," *Journal of Advanced Transportation*, vol. 2022, no. 1, p. 8281988, 2022.
- [31] W. Zhang, H. Liu, J. Han, Y. Ge, and H. Xiong, "Multi-agent graph convolutional reinforcement learning for dynamic electric vehicle charging pricing," in *Proceedings of the 28th ACM SIGKDD conference on knowledge discovery and data mining*, pp. 2471–2481, 2022.
- [32] X. Zhao, Y. Zhang, X. Zhang, Y. Yang, Y. Xie, Y. Li, and J. Luo, "Urban-focused multi-task offline reinforcement learning with contrastive data sharing," in *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 4512–4523, 2024.
- [33] H. Tan, Y. Yuan, H. Yan, S. Zhong, and Y. Yang, "Human preference-aware rebalancing and charging for shared electric micromobility vehicles," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2024.
- [34] A. Y. Lam, Y.-W. Leung, and X. Chu, "Autonomous-vehicle public transportation system: Scheduling and admission control," *IEEE Transactions on Intelligent Transportation Systems*, vol. 17, no. 5, pp. 1210–1226, 2016.
- [35] J. Ricketts, D. Barry, W. Guo, and J. Pelham, "A scoping literature review of natural language processing application to safety occurrence reports," *Safety*, vol. 9, no. 2, p. 22, 2023.
- [36] H. Zhang, R. Zhang, and J. Sun, "Developing real-time iot-based public safety alert and emergency response systems," *Scientific Reports*, vol. 15, no. 1, p. 29056, 2025.
- [37] N. A. B. . S. A. (NABSA), "2024 shared micromobility state of the industry report." <https://nabsa.net/about/industry/>, 2024.
- [38] O. International Transport Forum (ITF), "Safer micromobility: Technical background report," tech. rep., ITF, OECD, 2024.
- [39] P. de Souza Coelho, R. A. de Albuquerque Nobrega, L. K. de Oliveira, and M. Humberto, "The trade-off between equity and quality in public transportation: lessons from a brazilian case study," *npj Sustainable Mobility and Transport*, vol. 2, no. 1, p. 22, 2025.
- [40] T. F. Welch and S. Mishra, "A measure of equity for public transit accessibility," *Transportation Research Part A: Policy and Practice*, vol. 88, pp. 1–13, 2013.