

i-Rebalance: Personalized Vehicle Repositioning for Supply Demand Balance

Haoyang Chen¹, Peiyan Sun¹, Qiyuan Song¹, Wanyuan Wang¹,
Weiwei Wu¹, Wencan Zhang², Guanyu Gao³, Yan Lyu^{1*}

¹Southeast University, China

²National University of Singapore

³Nanjing University of Science and Technology

{hy_chen, spy, qiyuan_song, wywang, weiweiwu}@seu.edu.cn, wencanz@u.nus.edu, gygao@njust.edu.cn, lvyanly@seu.edu.cn

Abstract

Ride-hailing platforms have been facing the challenge of balancing demand and supply. Existing vehicle repositioning techniques often treat drivers as homogeneous agents and relocate them deterministically, assuming compliance with the repositioning. In this paper, we consider a more realistic and driver-centric scenario where drivers have unique cruising preferences and can decide whether to take the recommendation or not on their own. We propose i-Rebalance, a personalized vehicle repositioning technique with deep reinforcement learning (DRL). i-Rebalance estimates drivers' decisions on accepting repositioning recommendations through an on-field user study involving 99 real drivers. To optimize supply-demand balance and enhance preference satisfaction simultaneously, i-Rebalance has a sequential repositioning strategy with dual DRL agents: Grid Agent to determine the repositioning order of idle vehicles, and Vehicle Agent to provide personalized recommendations to each vehicle in the pre-defined order. This sequential learning strategy facilitates more effective policy training within a smaller action space compared to traditional joint-action methods. Evaluation of real-world trajectory data shows that i-Rebalance improves driver acceptance rate by 38.07% and total driver income by 9.97%. The code for our approach is available at <https://github.com/Haoyang-Chen/i-Rebalance>.

Introduction

Taxi and ride-hailing services have been facing supply-demand imbalances. Thousands of drivers spend almost 50% of their working time cruising on road for passengers (Zong et al. 2018). Meanwhile, thousands of travel demands got unserved due to limited vehicle supply nearby. There have been efforts on repositioning idle vehicles to demanding locations for supply-demand balance with deep reinforcement learning (DRL) techniques (Qin et al. 2022; Yang et al. 2020; Liu et al. 2021). By interacting with actual travel demands, these methods can adapt to the complex and dynamic environment and coordinate competing drivers.

However, these techniques hypothesize that *all the drivers unquestionably adhere to the repositioning recommendation and follow the recommended cruising route*. This is impractical because every driver has a unique cruising pref-




Figure 1: Impact of recommendation order. a) Initial scenario: Recommending four idle vehicles in Grid 5. Each prefers two neighboring grids, rejecting non-preferred options to stay at the current grid. The supply-demand gap δ is color-coded: red signifies shortage, and green indicates excess idle vehicles. b) Two recommendation orders $C \succ D \succ B \succ A$ (top row) and $A \succ B \succ C \succ D$ (bottom row) lead to different supply-demand balances due to drivers' diverse preferences.

erence (Zhang et al. 2020a,b). Imposing repositioning mandates on drivers would lead to unpleasant working experiences and cause driver loss on ride-hailing platforms.

To this end, we are motivated to consider a more realistic and driver-centric scenario where drivers could freely accept or reject repositioning recommendations. This requires generating personalized repositioning recommendation to improve the probability of acceptance. However, the challenges are:

1) *Subjective driver acceptance to repositioning recommendation*. Drivers can decide whether to accept recommendations in most ride-hailing platforms on their own. However, there has been lacking studies on under which circumstances and how likely the driver would accept the repositioning recommendation. Such decision-making from real drivers should be collected for accurate decision modeling.

2) *Personalized repositioning*. To incorporate unique preferences of drivers, we have to jointly decide each individual driver to visit a specific location. However, learning such a joint action can be challenging due to the large action space. For a smaller action space, we could reposition vehicles one by one, however, *the repositioning order of vehicles impacts the resulting supply-demand balance*, due to the fact that drivers with diverse preferences may reject unsatisfactory recommendation. Take Figure 1 as an example of recommending four vehicles in Grid 5 to address the supply-demand gap

*Corresponding author.

in neighboring grids. Each vehicle prefers two neighboring grids, rejecting non-preferred options to stay at the current grid. If we recommend in the order of $C \succ D \succ B \succ A$ (top row), drivers C and D will visit more demanding Grids 1 and 9 based on their preference. But the subsequent drivers B and A will reject recommendations to address the demand at Grid 2, leading to an imbalanced supply-demand. Conversely, the order $A \succ B \succ C \succ D$ (bottom row) achieves a more favorable outcome. All drivers visit preferred grids, simultaneously mitigating the supply-demand gap. Therefore, it is necessary to seek for such an order to help with better supply-demand balance and preference satisfaction.

In this paper, we propose i-Rebalance, a personalized vehicle repositioning technique to balance supply and demand. To estimate drivers' subjective decisions on reposition recommendation, we recruited 99 real drivers and conducted an on-field user study to collect their decision makings. For personalized repositioning, i-Rebalance sequentially repositions vehicles with dual DRL agents: Grid Agent and Vehicle Agent. Grid Agent learns *the repositioning order of idle vehicles* within a grid, while Vehicle Agent learns *personalized repositioning to each vehicle* in the pre-defined order. By joint training of the two agents, i-Rebalance can simultaneously optimize supply-demand balance and driver acceptance, facilitating policy training with a small action space. In summary, our contributions are:

- A personalized vehicle reposition technique i-Rebalance to balance supply and demand with DRL. It improves cruising preference satisfaction of drivers and their acceptance of repositions.
- A user study on 99 recruited ride-hailing drivers to estimate the likelihood of reposition acceptance.
- A sequential vehicle reposition framework with dual DRL agents to determine the reposition order and then the personalized reposition for each vehicle by the order.
- Evaluation on real-world taxi trajectory data shows that i-Rebalance greatly improves reposition acceptance of drivers and total driver income compared to baselines.

Related Work

Recently, deep reinforcement learning (DRL) has been widely used in vehicle repositioning problems (Qin et al. 2022; Liu et al. 2022; Jiao et al. 2021). By interacting with a simulated ride-hailing system, DRL techniques are able to capture the interactions between repositioning actions and dynamic demands and can coordinate competing vehicles. Centralized strategies have been explored, where a single agent is trained to govern reposition for all vehicles (He and Shin 2020; Oda and Joe-Wong 2018; Liu et al. 2020). However, such approaches suffer from extensive joint action search spaces. In contrast, decentralized approaches consider each vehicle as an independent agent (Al-Abbasi et al. 2019; Wen et al. 2017), but face the non-stationary exploration problems in multiagent RL tasks (Gao et al. 2022).

To bridge traditional vehicle repositioning methods (Wang et al. 2018; Xie et al. 2018; Xu et al. 2018) and DRL, some integrate request-matching and vehicle-repositioning into a single reinforcement learning framework (Xu et al.

2023), while others leverage hierarchical reinforcement learning to coordinate vehicles (Jin et al. 2019; Xi et al. 2022). Further efforts focus on simplifying the model, such as treating multiple agents in the same location as homogeneous agent to share the same policy network (Lin et al. 2018); averaging actions of agents in the same grids (Li et al. 2019); and combining vehicle-location maximum matching optimization during DRL training (Liu et al. 2021).

However, these methods assume each driver complies with the reposition recommendations unconditionally. In practice, drivers have their own preference on cruising routes (Zhang et al. 2020a,b) and can reject reposition recommendations in most ride-hailing platforms. Recently, (He and Shin 2020) includes collective driver preferences in their repositioning system that helps recommend locations that most drivers like to visit, without personalization. (Xu et al. 2020) integrates driver acceptance into its reposition model, but also overlooks personalized driver preferences. In contrast, i-Rebalance is the first to consider drivers' personalized cruising preferences and their freedom of acceptance of reposition. Specifically, we predict driver preferences with a LSTM network and conduct a user study to estimate the likelihood of acceptance. We propose a novel sequential vehicle reposition framework to learn reposition order of idle vehicles and to learn personalized reposition for each vehicle in the pre-defined order, respectively.

Problem Statement

We split the entire city into nonoverlapping rectangle grid map (He and Shin 2020), and time of a day into equal-length time intervals (e.g., 10 minutes). Based on this setting, we formally define the following concepts and the problem of personalized reposition recommendation.

Driver Cruising Preference. Each taxi driver has a unique preference for cruising routes for passengers. For example, some may prefer to hang around hot locations such as train stations and hospitals, some may prefer familiar locations near home. We define the cruising preference of driver i located in grid g at time t , denoted by ρ_t^i , by the probability of how likely the driver is willing to visit the neighboring 3×3 grids (including the current grid g). This will be estimated from the driver's historical cruising trajectories.

Reposition Recommendation. Ride-hailing platforms often provide recommendations on where to hunt for passengers to balance supply and demand. Instead of recommending an arbitrary location (Liu et al. 2022) on the map, we define a reposition recommendation at time t as visiting one of the neighboring 3×3 grids (including the current grid g) for smaller decision-making space (Liu et al. 2020).

Driver Acceptance towards reposition recommendation. We define a driver's acceptance as the probability of accepting the recommendation, given the driver's own cruising preference. This will be estimated by our on-field user study.

Supply-Demand Gap. For each grid g at time t , we define supply-demand gap, denoted by δ_t^g , as the difference between the number of idle vehicles and the number of ride requests within grid g at time t .

Problem Statement. i-Rebalance aims to generate a reposition recommendation for each idle vehicle at each




Figure 2: Overview of i-Rebalance. i-Rebalance comprises two phases: 1) *Driver Behavior Modeling* simulates realistic driver decision-making by predicting their cruising preferences and reposition acceptance probabilities. 2) *Sequential Vehicle Reposition with Dual DRL Agents* interacts with the simulator. Grid Agent observes nearby supply-demand gap Δ and driver preference ρ and determines the repositioning order of idle vehicles within the grid, e.g., $C \succ D \succ B \succ A$. By this order, Vehicle Agent observes individual preference ρ^i of driver i and real-time updated supply-demand gap Δ^i , and recommends reposition destinations for this vehicle. It receives rewards of supply-demand balance and preference satisfaction for each recommendation, while Grid Agent receives the average rewards after all recommendations.

time step t to minimize the overall absolute supply-demand gap and maximize the driver acceptance at the same time.

Next, we present the overview of i-Rebalance, followed by two driver behavior models to predict driver preferences and their acceptance of recommendations, and a sequential vehicle reposition framework with dual DRL agents.

System Overview

i-Rebalance comprises two distinct phases: 1) *Driver Behavior Modeling* for realistic driver decision-making simulation, and 2) *Sequential Vehicle Reposition with Dual DRL Agents* to learn personalized recommendation via interacting with the simulated drivers and environment (see Figure 2).

Driver Behavior Modeling predicts each driver’s cruising preferences based on their own historical trajectories, and estimates the probability of a driver accepting a reposition recommendation through an on-field user study. The two models are then integrated into the simulator, enabling us to simulate drivers’ cruising routes and their likelihood of accepting a reposition recommendation.

Sequential Vehicle Reposition with Dual DRL Agents interacts with the simulator to optimize vehicle repositioning across grids. To provide personalized reposition, we have to decide the assignment of each individual driver to a specific location. Training for such a joint-action faces challenges of a vast action space. A potential solution is to reposition vehicles one by one in a smaller action space. However, this introduces a new challenge: the repositioning order significantly impacts supply-demand balance. This is because drivers have diverse preferences and may reject unsatisfactory recommendation. Earlier driver decisions affect subsequent ones and thereby the resulting supply-demand balance (as illustrated in Figure 1).

Motivate by this, we seek for the optimal repositioning order that can help with better supply-demand balance and preference satisfaction. Specifically, we employ a Grid Agent to determine the reposition order first. By this order, a Vehicle Agent provides personalized recommendation for

each vehicle. Grid Agent observes nearby supply-demand gap and cruising preference of drivers within the grid while Vehicle Agent focuses on each vehicle’s own preference. Vehicle Agent receives intermediate reward for supply-demand balance and preference satisfaction, Grid Agent receives the average of Vehicle Agent’s reward for using its order. By training the two agents jointly, both supply-demand balance and preference satisfaction can be optimized simultaneously.

Driver Behavior Modeling

Driver Cruising Preference

We propose a lightweight LSTM based network, to predict driver’s next cruising directions based on individual past cruising trajectories. The prediction will be used to 1) estimate driver acceptance to a recommendation, and 2) simulate preferred cruising directions upon rejection.

As illustrated in Figure 3, the two-layer LSTM network takes the input of factors including i) the driver’s own working states, such as how long she has been cruising and how much she has earned, and ii) the current location, which can be presented by distances to a set of POIs, traffics, supply and demand nearby. It predicts probabilities ρ_t of the driver willing to visit the neighboring 3×3 grids. See the detailed feature engineering and network description in Appendix A. This lightweight network improves predictions for drivers with limited historical cruising data, and it’s more efficient than larger neural networks (Zhang et al. 2020b,a). However, existing trajectory prediction networks are also applicable.

Driver Acceptance to Reposition Recommendation

To estimate the likelihood of a driver accepting a reposition recommendation, we conducted an on-field user study to collect drivers’ decisions in various working scenarios.

We recruited drivers from the Didi platform by requesting rides at a few selected origin locations (including business centers, hospitals, universities, and scenic spots) to a random destination in our city. After a driver comes for the pick-up, she is asked to park her vehicle at a safe place (mostly at

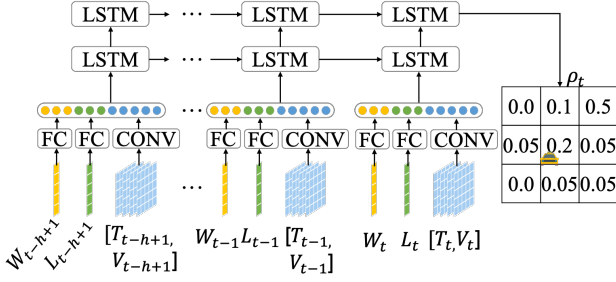


Figure 3: Cruising preference prediction network. The network takes a sequence of the features including driver states W_t , POI distances L_t , traffics T_t , and supply-demand features V_t from time $t - h + 1$ to t as input and predicts probabilities ρ_t of the driver visiting the neighboring 3×3 grids.

the origin) for an interview. Drivers followed the procedure: 1) The driver answers demographic questions such as working experience and income per ride request. She is also asked to estimate how often she accepts the reposition instructions from the platform (Obedience). 2) She proceeds to the main survey with 3 trials of different maps with different decision-making questions. For each trial, she sees a map with 5×5 grids of a region in the city and is asked whether she is familiar with the region. She is then asked to imagine herself located in the center grid with an empty vehicle and to rank the neighboring 9 grids (including the center grid) as the next preferred cruising direction. We then provide a repositioning recommendation, i.e., visit one of the 9 ranked grids (Preference Ranking), as well as Expected Income if taking the recommendation, and ask the driver to decide whether to accept the recommendation or not. For every trial, each driver answers 10 decision-making questions, each with a randomly selected ranked grid as a recommendation and one of 7 expected incomes ranging from 6 to 16.

We recruited 99 Didi drivers at 11 locations, with working experience 0.1-8 (Mean = 3.2) years, and reported income of 10 - 20 (Mean = 14.6) RMB per ride. We excluded 6/105 participants who did not complete all questions or gave unreliable answers. Participants completed the survey questions within 20 minutes and were compensated 12 RMB (\$1.50 USD). We formulated whether a driver would accept or reject our reposition recommendation as a binary classification problem. We fit a binomial logistic regression model with *Preference Ranking* (denoted as r), *Expected Income* (denoted as m), and *Driver's Obedience* (denoted as o) as independent variables:

$$P_{accept} = 1/(1 + \exp(-(b + w_r r + w_m m + w_o o))). \quad (1)$$

Variable	Coef.	SE	χ^2	p	R2
Intercept	-1.31	0.26	25.06	<.0001	
Preference Ranking	-0.44	0.03	273.50	<.0001	.3794
Expected Income	0.29	0.02	172.75	<.0001	
Driver's Obedience	2.17	0.22	95.93	<.0001	

Table 1: Statistical analysis of driver responses on the binomial logistic regression model predicting the likelihood to accept a repositioning recommendation.




Figure 4: Partial dependence plot of independent variables.

We found that all three factors significantly influenced drivers' willingness to accept the recommendation (Table 1). The model achieved a classification accuracy of 76.37% and an area under the ROC curve (AUC) of 0.8266. According to the partial dependence plot (Figure 4), we also observed that drivers have a higher chance of rejecting the less preferred recommendation (lower ranked). This indicates that they are more willing to collaborate with the system if it can provide highly personalized recommendations. Conversely, drivers are less likely to reject the recommendation if the recommendation helps to bring a higher expected income. Also, a driver who is inherently more obedient exhibits a higher likelihood of accepting the recommendation.

These statistical results and analysis indicate a good fit of the binomial logistic regression model. We use this model (Eq. 1) with $b = -1.31$, $w_r = -0.44$, $w_m = 0.29$, $w_o = 2.17$ (see Table 1) to simulate driver acceptance to a reposition recommendation later.

Sequential Reposition with Dual DRL Agents

In this section, we present a sequential vehicle reposition framework with dual DRL agents to learn personalized reposition recommendations.

Order Determination by Grid Agent

We define a **Grid Agent** to learn the optimal order for recommending idle vehicles within a grid. Its state, action, and reward are defined as follows:

State of Grid Agent, denoted by S_G , includes the predicted supply-demand gap in neighboring grids at time step $t + 1$, and preferences of available drivers in this grid, i.e.,

$$S_G = \langle \Delta_{t+1}, \rho_t^1, \rho_t^2, \dots, \rho_t^n \rangle, \quad (2)$$

where Δ_{t+1} denotes a vector of supply-demand gaps of neighboring 3×3 grids (including the current grid) at time $t + 1$, i.e., $\Delta_{t+1} = (\delta_{t+1}^1, \delta_{t+1}^2, \dots, \delta_{t+1}^9)$. The supply-demand gap δ_{t+1}^k of grid k is estimated from the number of idle vehicles (supply) that will stay or arrive at the grid, subtracted by the number of requests (demand) at time $t + 1$ predicted by a recurrent graph convolution network (Wang et al. 2020). ρ_t^i denotes the preference (i.e., the probability of visiting each of the 9 grids) of i -th driver, $1 \leq i \leq n$. Note that the number of available drivers n in a grid is dynamic. To ensure dimension consistency for state features, we employ a preference matrix capable of accommodating a larger driver count with non-driver entities padded with zeros.

Action of Grid Agent, denoted by A_G , is a vector of priority scores assigned to each driver i , i.e., $A_G = (a_1, a_2, \dots, a_n)$, where $0 \leq a_i \leq 1$, $1 \leq i \leq n$. Later

on, Vehicle Agent will recommend the drivers based on the order $\succ = (\succ_1, \dots, \succ_n)$ ranked by their priority scores.

Reward of Grid Agent, denoted by R_G , is the average reward received by Vehicle Agent after all the idle vehicles get recommendations and repositioned by the output order.

Personalized Recommendation by Vehicle Agent

We define a **Vehicle Agent** to learn personalized recommendations for repositioning idle vehicles within a grid, following the order determined by the Grid Agent. The definitions of state, action, and reward are detailed below.

State of a vehicle agent for recommending to driver i , denoted by S_V^i , includes the real-time updated supply-demand gap in future time $t + 1$, denoted by Δ_{t+1}^i , along with the cruising preference of driver i , denoted by ρ_t^i , i.e.,

$$S_V^i = \langle \Delta_{t+1}^i, \rho_t^i \rangle, \quad (3)$$

Δ_{t+1}^i is calculated from the initial supply-demand gap Δ_{t+1} before recommending for the grid and the reposition decisions made by drivers leading up to the i -th driver's turn. ρ_t^i denotes the preference of driver i .

Action of a vehicle agent, denoted by a_V^i , is to recommend driver i to visit one of the 9 grids g . Different from deterministic repositioning methods, our action simply provides a reposition recommendation, whether the driver accepts the recommendation or insists on her own cruising preferences is determined by herself.

Reward of a recommendation action should be able to measure supply-demand balance, as well as preference satisfaction to the recommendation. Specifically,

Balance Reward, denoted by $R_B^i(S_V^i, a_V^i)$, is a standardized supply-demand gap at the recommended grid a_V^i , i.e.,

$$R_B^i(S_V^i, a_V^i) = -(\delta_{t+1}^{a_V^i} - \mu_g) / \sigma_g, \quad (4)$$

Here, μ_g and σ_g denote the mean and standard deviation of the supply-demand gap of the neighboring 9 grids of the grid g , where g is the current location of driver i . This reward guides drivers to visit the most demanding grids.

Preference Satisfaction Reward, denoted by $R_P^i(S_V^i, a_V^i)$, measures how much a reposition recommendation satisfies driver's preference. It is the min-max normalized rankings of driver i 's preferences on the 9 neighboring grids. A higher-ranking recommendation gets a higher preference reward.

Overall, the total reward of a reposition recommendation is the summation of the two types of rewards, i.e.,

$$R_V^i(S_V^i, a_V^i) = \alpha_B R_B^i(S_V^i, a_V^i) + \alpha_P R_P^i(S_V^i, a_V^i), \quad (5)$$

where α_B and α_P are weights of corresponding rewards.

Reposition Algorithm

i-Rebalance adopts an A2C network to learn personalized reposition policies. Algorithm 1 depicts general procedures of training i-Rebalance. We define an episode as a day and divide the daytime into equal-length time steps.

The algorithm initializes Grid Agent network θ_G and Vehicle Agent network θ_V with vacant memory buffers M_G and M_V , respectively. At time step t , we predict driver preference by our LSTM network (Fig. 3) and the demand in

Algorithm 1: Personalized Taxi Repositioning

Input: Max episodes N , total time step T

Output: Grid Agent θ_G , Vehicle Agent θ_V

```

1: Memory  $M_V, M_G \leftarrow \emptyset$ , randomly initialize network  $\theta_V, \theta_G$ .
2: while episode  $n < N$  do
3:   Simulator reset
4:   while  $t < T$  do
5:     Predict demands at  $t + 1$  and driver preference
6:     for each grid in radial sequence do
7:       Get state  $S_G$  for grid agent
8:       Generate a repositioning order  $A_G$  by network  $\theta_G$ 
9:       Store  $[S_G, A_G]$  to  $M_G$ 
10:      for each available driver following  $A_G$  do
11:        Get state  $S_V$  for vehicle agent
12:        Recommend a reposition  $A_V$  by network  $\theta_V$ 
13:        Driver decides to execute  $A_V$  or visit preferred grid
14:        Get rewards  $R_B$  and  $R_P$ 
15:        Store  $[S_V, A_V, R_P + R_B]$  to  $M_V$ 
16:      end for
17:      Get average reward  $R_G$  and update to  $M_G$ 
18:    end for
19:    Updates real demand and supply
20:    Match demands to their nearest idle vehicles
21:    Update next state  $S_V'$  to  $M_V$  for repositioned vehicles
22:    Replay from buffer  $M_V$  to update  $\theta_V$ 
23:    Replay from buffer  $M_G$  to update  $\theta_G$ 
24:  end while
25: end while

```

each grid at time $t + 1$. Then, we start to reposition idle vehicles in each grid in a radial order starting from the center grid and spiraling outward (Lines 6 to 18). The Grid Agent θ_G observes the state S_G for a grid, generates idle vehicle order A_G , and stores this experience $[S_G, A_G]$ into memory M_G . Vehicle Agent then generates a reposition recommendation A_V for each idle vehicle order A_G . The simulator emulates driver's decision on whether to accept the recommendation or visit the preferred grid by our binomial logistic regression model (Eq. 1). Vehicle Agent gets immediate Reward R_V^i after recommending vehicle i . The transition $[S_V, A_V, R_V^i]$ is stored to replay buffer M_V . Once all idle vehicles in the grid are repositioned, Grid Agent gets the average reward R_G of all those immediate rewards. After all idle vehicles are processed, our simulator updates demands and vehicle supplies by repositions outcome and the drop-offs of occupied vehicles (Line 19) for time $t + 1$. The demands are then matched to their nearest idle vehicles. Finally, replay buffers are updated by the next state and both Grid Agent and Vehicle Agent are updated by the newest buffers separately.

We further prove that our sequential reposition framework can achieve the same total rewards as learning joint action of repositioning all idle vehicles simultaneously in a grid.

Theorem 1. *Given the state at the grid g , $S_G = \langle \Delta_{t+1}, \rho_t^1, \rho_t^2, \dots, \rho_t^n \rangle$, the proposed sequential reposition framework learns reposition order and reposition policy can achieve the same reward with that of the optimal policy of assigning vehicles to locations simultaneously.*

Proof. At the state S_G , assume that there are n idle vehicles being repositioned and the optimal joint reposition action

$\vec{a}^* = \langle a_1^*, a_2^*, \dots, a_n^* \rangle$, which has the expected immediate global reward $\mathbb{E}_{P_{accept}}[\sum_{r \in Neg(g)} \alpha_B R_B^r + \sum_{i=1}^n \alpha_P R_P^i]$, where $Neg(g)$ denote the neighboring grids of g and $R_B^r = -(\delta_{t+1}^r - \mu_g)/\sigma_g$ denotes the balance reward of the neighboring grid r . Given the joint action \vec{a}^* , and a randomized order $\succ = (\succ_1, \dots, \succ_n)$, then the same immediate expected global reward $\mathbb{E}_{P_{accept}}[\sum_{r=1}^9 \alpha_B R_B^r + \sum_{i=1}^n \alpha_P R_P^i]$ will be achieved by the Grid agent if each vehicle i is repositioned by the action a_i^* in sequence according to the order \succ . In our sequential reposition mechanism, the reposition action of the latter vehicle agent depends on the former vehicle agent’s action. Thus, given the $(j - 1)$ th vehicle’s action, if the j th vehicle agent would like to change his action from a_j^* to a_j' , a larger individual reward must be achieved, i.e.,

$$R_V^j(S_V^j, a_j') = \alpha_B R_B^i(S_V^j, a_j') + \alpha_P R_P^j(S_V^j, a_j') > R_V^j(S_V^j, a_j^*).$$

Because the immediate global reward received by the joint action (i.e., $\mathbb{E}_{P_{accept}}[\sum_{r=1}^9 \alpha_B R_B^r + \sum_{i=1}^n \alpha_P R_P^i]$) is the sum of the individual vehicle agent’s reward, improving the action from a_j^* to a_j' also improves the immediate global reward. In summary, for any joint reposition policy $\vec{a}^* = \langle a_1^*, a_2^*, \dots, a_n^* \rangle$, our sequential reposition strategy can also find a reposition order and reposition actions that produce immediate global reward not less than that of \vec{a}^* . \square

Evaluation

In this section, we evaluate the performance of i-Rebalance in a simulator built with a real-world taxi trajectory dataset. In the following, we first describe experimental settings and then discuss experimental results.

Experimental Settings

Environment Simulator We employ and extend an open-source simulator (Liu et al. 2020) that simulates ride-hailing platform operations to receive ride requests and dispatch the requests to the nearest idle vehicles. The ride requests are sampled from a real-world taxi trajectory dataset, and driver starting times and locations are initialized accordingly. The simulator plans routes with Dijkstra shortest path algorithm and calculates travel time by estimating speed from the dataset for more realistic simulation.

We divided the city map into $1.2\text{km} \times 1.2\text{km}$ rectangle grids, and the day into 10 minutes time intervals (suitable for most vehicles to traverse a grid). The simulator predicts cruising directions to neighboring grids (or stay in the same grids) for each driver every time step by our LSTM network (Fig. 3). Meanwhile, the simulator emulates the driver decision makings on accepting reposition recommendations by Eq. 1. The driver either visits the recommended location, or randomly visits one of the top preferred four grids (See Appendix A for preference prediction performance).

Datasets We use a real-world taxi trajectory dataset collected in Chengdu, China, with a total of 14,865 taxis and 10,710,949 passenger trips within a month (Lyu et al. 2019). For experiment efficiency, we sampled 500 taxis and 243,679 passenger demands. The road network of Chengdu is obtained from OpenStreetMap, spanning the city center within 30.62 - 30.69 latitude and 107.01 - 107.11 longitude.

Baseline Techniques We compare the following methods:

- *No Reposition* does not reposition drivers but lets them cruise as estimated preferences in the simulator.
- *Random* recommends idle vehicles to a random one of neighboring 3×3 grids in a random order.
- *Demand Greedy* recommends vehicles to the current most demanding neighboring grid in a random order.
- *Reward Greedy* suggests the most demanding and satisfactory neighboring grid to maximize reward by Eq.5.
- *Min Cost Flow* (Ming et al. 2020) formulates vehicle repositioning as a min cost flow problem and minimizes repositioning cost while maximizing the total served demands.
- *Contextual DDQN* (Liu et al. 2020) uses DDQN to reposition idle vehicles and coordinates them by encoding prior driver decisions into subsequent driver observations.
- *Proportional Reposition* defines its action as redistributing proportions of idle vehicles from a grid to its neighbors (Mao et al. 2020). This facilitates coordination among idle vehicles akin to joint action methods (Liu et al. 2022), yet maintains a relatively smaller action space.
- *By Collective Preference* considers collective preference of drivers by their historical visiting frequency (He and Shin 2020), and thereby tends to recommend demanding grids that most drivers like to visit.

Metrics We evaluate the performance of all the methods with the following metrics.

- *Total Driver Income (TDI)* is the total income from all the served demands in a day. Better supply-demand balance help to achieve a higher TDI.
- *Reposition Income (RI)* is the total income earned from successful reposition recommendations. We use RI/TDI to measure how much is earned from repositions.
- *Request Response Rate (RRR)* is the ratio between the number of orders got served to the total number of orders.
- *Acceptance Rate* is the ratio of accepted repositions to total recommendations, indicating driver satisfaction.
- *Number of Repositions* is the number of accepted recommendations, excluding those for staying in the same grid. This assesses reposition cost for supply-demand balance.

Implementation In i-Rebalance, we implemented the actor and critic networks with two fully connected layers (64 units, *ReLU* activations), respectively. Grid Agent and Vehicle Agents use the same architecture, normalizing outputs with softmax and sigmoid functions separately. We used entropy loss and batch normalization, entropy- β is set to 0.01 and the batch size is 10. We set a decay factor $\gamma = 0.98$ and train the network with *Adam* optimizer (Liu et al. 2020). In reward function (Eq. 5), the weights $\alpha_B = 2$, and $\alpha_P = 1$ achieve the best performance. The models are trained on Intel Core i9-10940K CPU @3.30GHz, NVIDIA GeForce RTX 3090, and 32GB memory.

Experimental Results

Performance Comparison Table 2 summarizes the comparison results of our i-Rebalance against the baselines.

Methods	Norm. TDI	RI/TDI	RRR	Accept. Rate	#Repos.
No Reposition	100.00%	N/A	72.15%	N/A	10,733
Random	103.02%	88.85%	73.60%	48.39%	9,691
Demand Greedy	101.81%	85.73%	73.11%	45.44%	11,061
Reward Greedy	102.03%	87.36%	73.27%	50.50%	10828
Min Cost Flow	103.00%	87.48%	73.46%	44.68%	7,700
Contextual DDQN	102.46%	87.44%	73.33%	46.07%	10,866
Proportional Repositioning	102.59%	88.20%	73.60%	22.96%	6,749
By Collective Preference	109.00%	95.28%	77.29%	66.48%	6,803
i-Rebalance	109.97%	98.70%	78.35%	84.43%	6,625

Table 2: Performance Comparison of Vehicle Reposition Methods. Total Driver Income (TDI) is normalized by dividing that of *No Reposition*.

Overall, i-Rebalance performs the best in all metrics, indicating that i-Rebalance can achieve a good balance of supply and demand, providing the most profitable, satisfactory, and cost-effective reposition recommendations.

No Reposition has the lowest TDI by letting the drivers cruise by themselves without cooperating with each other. *Demand Greedy* has similar performance as it repositions vehicles to current most demanding grids, vehicles are easily gathered in some grids. *Reward Greedy* performs better than *Demand Greedy* by incorporating driver preference. *Random* has slightly higher TDI and ORR than *Reward Greedy* because it disperses vehicles everywhere. *Min Cost Flow* slightly worse than *Random* as it minimizes reposition cost. Overall, these non-learning-based methods have worse performance compared to the DRL methods, due to the naive strategies and short-term focus.

Among the DRL methods, *Contextual DDQN* has the worst performance as it coordinates vehicles by encoding early repositioning outcomes for state of later ones, which might not be sufficient. It also employs fixed rewards, less adaptable to an uncertain environment where drivers can reject repositions. *Proportional Repositioning* improves coordination via homogeneous joint action, yet ignores individual preference. *By Collective Preference* enhances performance with driver collective preference, achieving the highest acceptance rate (66.48%) among the baselines. But it still fails to consider personalized preference. In contrast, i-Rebalance incorporates personalized preference and coordinates drivers by learning their reposition order, achieving the highest driver income and acceptance rate.

Ablation Study We conducted two ablation studies to evaluate 1) effectiveness of learning recommendation order, and 2) effectiveness of incorporating personalized preference with different designs of state and reward.

Effectiveness of order learning. We design 3 variations of i-Rebalance. *Fixed-Order* employs a fixed order prioritizing drivers based on their preference distribution’s correlation with supply-demand gaps (measured by Pearson correlation). *Random-Order* employs a random order. *Joint-Action* involves a single agent to jointly determine

Methods	Norm. TDI	RI/TDI	Accept. Rate
Fixed-Order	109.89%	98.50%	83.40%
Random-Order	109.77%	98.44%	83.31%
Joint-Action	103.15%	90.18%	52.53%
i-Rebalance	109.97%	98.70%	84.43%

Table 3: Performance of i-Rebalance variations with different reposition order.

Methods	Norm. TDI	RI/TDI	Accept. Rate
RNP-SNP	109.37%	95.69%	67.95%
RP-SNP	109.40%	95.75%	68.11%
RNP-SP	109.73%	98.29%	81.65%
RP-SP	109.97%	98.70%	84.43%

Table 4: Performance of i-Rebalance variations with different state and reward design for Vehicle Agent. R: reward, S: state, P: preference, NP: no preference.

driver assignments to specific locations. Table 3 summarizes the comparison results. i-Rebalance excels by collaborating with the Vehicle Agent for the optimal order learning. *Fix-Order* is slightly less effective than i-Rebalance but surpasses *Random-Order* due to driver preference consideration. *Joint-Action* performs worst due to its larger action space, challenging to find the optimal solutions.

Effectiveness State and Reward Design for Personalization. We create 3 i-Rebalance variations with different state and reward structures for driver preference in Vehicle Agent. *RNP-SNP* lacks both preference satisfaction reward and preference observation in Vehicle Agent, *RNP-SP* has preference observation but not preference reward, and *RP-SNP* incorporates preference reward but lacks preference observation. i-Rebalance combines both reward and observation, known as *RP-SP*. For equitable comparison, we maintain the reposition order learned in i-Rebalance. Table 4 summarizes comparison results. The four methods display similar total driver income (TDI) and minor variations in RI/TDI, yet show significant differences in acceptance rates. This shows that i-Rebalance can generate more satisfactory recommendations without sacrificing total driver income. Also, integrating driver preferences in state (*RNP-SP*) is more effective than integrating that in reward (*RP-SNP*).

Conclusion

In this paper, we present i-Rebalance, a personalized vehicle reposition technique to balance supply and demand and to improve driver acceptance of reposition recommendations. i-Rebalance models driver cruising preferences with a lightweight LSTM network and estimates their decisions on accepting the recommendations through a user study. It sequentially repositions vehicles with dual DRL agents to learn the reposition order and personalized recommendation, respectively. Evaluation shows i-Rebalance greatly improves reposition acceptance of drivers and total driver income compared to all baselines. Future work includes generalizing our sequential learning strategy to coordinate multi-agents for personalized decision making, and exploring personalized reposition incentives for drivers.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China under Grant No. 62232004, No.61972086, No. 62102082, No. 62072099, No.61932007; in part by the Jiangsu Natural Science Foundation of China under Grant No. BK20210203, BK20230024.

References

- Al-Abbasi, A. O.; Ghosh, A.; and Aggarwal, V. 2019. Deep-pool: Distributed model-free algorithm for ride-sharing using deep reinforcement learning. *IEEE Transactions on Intelligent Transportation Systems*, 20(12): 4714–4727.
- Gao, G.; Wen, Y.; and Tao, D. 2022. Distributed energy trading and scheduling among microgrids via multiagent reinforcement learning. *IEEE Transactions on Neural Networks and Learning Systems*, 34: 10638 – 10652.
- He, S.; and Shin, K. G. 2020. Spatio-temporal capsule-based reinforcement learning for mobility-on-demand coordination. *IEEE Transactions on Knowledge and Data Engineering*, 34(3): 1446–1461.
- Jiao, Y.; Tang, X.; Qin, Z. T.; Li, S.; Zhang, F.; Zhu, H.; and Ye, J. 2021. Real-world ride-hailing vehicle repositioning using deep reinforcement learning. *Transportation Research Part C: Emerging Technologies*, 130: 103289.
- Jin, J.; Zhou, M.; Zhang, W.; Li, M.; Guo, Z.; Qin, Z.; Jiao, Y.; Tang, X.; Wang, C.; Wang, J.; et al. 2019. Coride: joint order dispatching and fleet management for multi-scale ride-hailing platforms. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, 1983–1992.
- Li, M.; Qin, Z.; Jiao, Y.; Yang, Y.; Wang, J.; Wang, C.; Wu, G.; and Ye, J. 2019. Efficient ridesharing order dispatching with mean field multi-agent reinforcement learning. In *The World Wide Web Conference*, 983–994.
- Lin, K.; Zhao, R.; Xu, Z.; and Zhou, J. 2018. Efficient large-scale fleet management via multi-agent deep reinforcement learning. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 1774–1783.
- Liu, C.; Chen, C.-X.; and Chen, C. 2021. Meta: A city-wide taxi repositioning framework based on multi-agent reinforcement learning. *IEEE Transactions on Intelligent Transportation Systems*, 23(8): 13890–13895.
- Liu, Y.; Wu, F.; Lyu, C.; Li, S.; Ye, J.; and Qu, X. 2022. Deep dispatching: A deep reinforcement learning approach for vehicle dispatching on online ride-hailing platform. *Transportation Research Part E: Logistics and Transportation Review*, 161: 102694.
- Liu, Z.; Li, J.; and Wu, K. 2020. Context-aware taxi dispatching at city-scale using deep reinforcement learning. *IEEE Transactions on Intelligent Transportation Systems*, 23(3): 1996–2009.
- Lyu, Y.; Liu, X.; Chen, H.; Mangal, A.; Liu, K.; Chen, C.; and Lim, B. 2019. OD morphing: Balancing simplicity with faithfulness for OD bundling. *IEEE Transactions on Visualization and Computer Graphics*, 26(1): 811–821.
- Mao, C.; Liu, Y.; and Shen, Z.-J. M. 2020. Dispatch of autonomous vehicles for taxi services: A deep reinforcement learning approach. *Transportation Research Part C: Emerging Technologies*, 115: 102626.
- Ming, L.; Hu, Q.; Dong, M.; and Zheng, B. 2020. An effective fleet management strategy for collaborative spatio-temporal searching: GIS cup. In *Proceedings of the 28th International Conference on Advances in Geographic Information Systems*, 651–654.
- Oda, T.; and Joe-Wong, C. 2018. MOVI: A model-free approach to dynamic fleet management. In *IEEE INFOCOM 2018-IEEE Conference on Computer Communications*, 2708–2716.
- Qin, Z. T.; Zhu, H.; and Ye, J. 2022. Reinforcement learning for ridesharing: An extended survey. *Transportation Research Part C: Emerging Technologies*, 144: 103852.
- Wang, R.; Chow, C.; Lyu, Y.; Lee, V. S.; Kwong, S.; Li, Y.; and Zeng, J. 2018. TaxiRec: Recommending Road Clusters to Taxi Drivers Using Ranking-Based Extreme Learning Machines. *IEEE Transactions on Knowledge and Data Engineering*, 30(03): 585–598.
- Wang, X.; Ma, Y.; Wang, Y.; Jin, W.; Wang, X.; Tang, J.; Jia, C.; and Yu, J. 2020. Traffic flow prediction via spatial temporal graph neural network. In *Proceedings of The Web Conference 2020*, 1082–1092.
- Wen, J.; Zhao, J.; and Jaillet, P. 2017. Rebalancing shared mobility-on-demand systems: A reinforcement learning approach. In *2017 IEEE 20th International Conference on Intelligent Transportation Systems*, 220–225.
- Xi, J.; Zhu, F.; Ye, P.; Lv, Y.; Tang, H.; and Wang, F.-Y. 2022. Hmdrl: Hierarchical mixed deep reinforcement learning to balance vehicle supply and demand. *IEEE Transactions On Intelligent Transportation Systems*, 23(11): 21861–21872.
- Xie, X.; Zhang, F.; and Zhang, D. 2018. PrivateHunt: Multi-source data-driven dispatching in for-hire vehicle systems. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2(1): 1–26.
- Xu, J.; Rahmatizadeh, R.; Bölöni, L.; and Turgut, D. 2018. Taxi dispatch planning via demand and destination modeling. In *2018 IEEE 43rd Conference on Local Computer Networks*, 377–384.
- Xu, M.; Yue, P.; Yu, F.; Yang, C.; Zhang, M.; Li, S.; and Li, H. 2023. Multi-agent reinforcement learning to unify order-matching and vehicle-repositioning in ride-hailing services. *International Journal of Geographical Information Science*, 37(2): 380–402.
- Xu, Z.; Men, C.; Li, P.; Jin, B.; Li, G.; Yang, Y.; Liu, C.; Wang, B.; and Qie, X. 2020. When recommender systems meet fleet management: Practical study in online driver repositioning system. In *Proceedings of The Web Conference 2020*, 2220–2229.
- Yang, Y.; Wang, X.; Xu, Y.; and Huang, Q. 2020. Multiagent reinforcement learning-based taxi predispatching model to balance taxi supply and demand. *Journal of Advanced Transportation*, 2020: 1–12.

Zhang, X.; Li, Y.; Zhou, X.; and Luo, J. 2020a. cgail: Conditional generative adversarial imitation learning—an application in taxi drivers’ strategy learning. *IEEE transactions on big data*, 8(5): 1288–1300.

Zhang, X.; Li, Y.; Zhou, X.; Zhang, Z.; and Luo, J. 2020b. Trajgail: Trajectory generative adversarial imitation learning for long-term decision analysis. In *2020 IEEE International Conference on Data Mining (ICDM)*, 801–810.

Zong, F.; Wu, T.; and Jia, H. 2018. Taxi drivers’ cruising patterns—Insights from taxi GPS traces. *IEEE Transactions on Intelligent Transportation Systems*, 20(2): 571–582.