# Data-driven optimization of pricing and vehicle relocation for ridesourcing platforms considering reservation

Xiaoshu Ding[a], Haoning Xi[b], Han Fang[a], and Sisi Jian[a,*]

[a] Department of Civil and Environmental Engineering, The Hong Kong University of Science and Technology, Hong Kong

[b] Newcastle Business School, The University of Newcastle, Sydney, Australia

[*] Corresponding author: cesjian@ust.hk

*Extended abstract submitted for presentation at the 12th Triennial Symposium on Transportation Analysis conference (TRISTAN XII)*
*June 22-27, 2025, Okinawa, Japan*

**Keywords:** Ridesourcing, Data-driven optimization, Long Short-Term Memory (LSTM), Hybrid offline-online framework, Deep Reinforcement Learning

## 1 INTRODUCTION

In response to the growing demand for pre-booked trips, ridesourcing platforms, such as Uber and Didi Chuxing, have introduced reservation services, aiming to ensure punctual departures for passengers while reducing idle time for drivers. However, many customers have voiced frustrations over the trip fares, which can be several times higher than real-time orders. Such disparity raises concerns over fairness and trust in the pricing strategies of these ridesourcing platforms. Ideally, the pricing of pre-booked orders should reflect the demand-supply dynamics at the actual departure time, which are, however, not revealed in advance. Traditional dynamic pricing raises prices when demand exceeds supply to maintain balance. In contrast, pre-booked orders require guaranteed fulfilment, making it challenging to apply such price adjustments. Instant price hikes could lead to unexpectedly high costs for travellers or unfulfilled orders, thereby affecting customer satisfaction and retention. Therefore, it is essential to design a transparent and fair pricing scheme for reservation services. Furthermore, pricing decisions impact vehicle relocation, which is critical to efficiently meeting travellers' requests and optimizing mobility resource utilization. While existing research has explored vehicle dispatching and relocating optimization for ridesourcing Yahia *et al.* (2021) and ride-pooling services Engelhardt *et al.* (2022) considering the difference between pre-booked and on-demand customers, few studies focus on the interaction between pricing schemes for pre-booked and on-demand services and the joint pricing and relocation strategies for ridesourcing platforms.

To fill this gap, we propose a data-driven framework that combines partial offline optimization model for pre-booked orders with adaptive online optimization model for real-time orders and vehicle relocation strategies. We first employ deep learning neural networks (i.e., Long short-term memory) to predict demand functions based on the historical data. These predictions are then incorporated into a profit-maximization problem to determine the optimal pricing strategies for pre-booked orders, which can proactively balance supply and demand. However, since real demand may still deviate from these predictions due to uncertainties, adaptive vehicle relocation and dynamic pricing strategies are required for real-time orders. To handle these fluctuations, we use Deep Reinforcement Learning (DRL)(Huang *et al.*, 2022) to adjust pricing, dispatching, and relocation decisions based on real-time demand, while maintaining pre-booked commitments.

## 2 METHODOLOGY

We formulate the offline pricing and vehicle allocation problem for a ridesourcing platform operating in a zonal transport network. Let $\mathcal{N} = \{1, 2, ..., N\}$ represent the set of zones and

$\mathcal{T} = \{1, 2, ..., T\}$ the time slots. In each time slot $t$, the platform sets prices for both pre-booked and real-time orders, denoted by $p_{ij,t}^{PB}$ and $p_{ij,t}^{RT}$, respectively. The demand for these services depends on passengers' willingness-to-pay, which is modeled using a distribution function $f$. The number of passengers willing to pay for trips from zone $i$ to zone $j$ is given by $n_{ij,t}^{PB}(p_{ij,t}^{PB})$ and $n_{ij,t}^{RT}(p_{ij,t}^{RT})$, influenced by the potential demand $D_{ij,t}^{PB}$ and $D_{ij,t}^{RT}$. The platform has full knowledge of these demand distributions and adjusts pricing accordingly. While pre-booked orders are guaranteed, real-time orders may face unmet demand due to vehicle shortages. Therefore, the platform decides the number of vehicles to allocate for real-time orders, $v_{ij,t}^{RT}$, to meet the highest demand. The number of successfully fulfilled pre-booked and real-time orders, $O_{ij,t}^{PB}$ and $O_{ij,t}^{RT}$, are then determined. The offline profit maximization problem is then formulated as follows:

$$\max_{\boldsymbol{p},\boldsymbol{d},\boldsymbol{x}} \sum_{i,j} \sum_{t} \left(p_{ij,t}^{RT} O_{ij,t}^{RT} + p_{ij,t}^{PB} O_{ij,t}^{PB} - c_{ij} x_{ij,t}\right) \tag{1}$$

s.t.

$$\sum_{j \in \mathcal{N}, j \neq i} \left(v_{ij,t}^{RT} + O_{ij,t}^{PB} + x_{ij,t}\right) \leqslant s_{i,t}, \quad \forall i \in \mathcal{N}, t \in \mathcal{T} \tag{2}$$

$$O_{ij,t}^{RT} = \min\left\{n_{ij,t}^{RT}(p_{ij,t}^{RT}), v_{ij,t}^{RT}\right\}, \quad \forall i, j \in \mathcal{N}, i \neq j, t \in \mathcal{T} \tag{3}$$

$$O_{ij,t}^{PB} = n_{ij,t}^{PB}(p_{ij,t}^{PB}), \quad \forall i, j \in \mathcal{N}, i \neq j, t \in \mathcal{T} \tag{4}$$

$$\begin{aligned} q_{i,t+1} = q_{i,t} &- \sum_{j \in \mathcal{N}, j \neq i} \left(v_{ij,t}^{RT} + O_{ij,t}^{PB} + x_{ij,t}\right) \\ &+ \sum_{j \in \mathcal{N}, j \neq i} \left(v_{ji,t-\delta_{ji}}^{RT} + O_{ji,t-\delta_{ji}}^{PB} + x_{ji,t-1}\right), \quad \forall i \in \mathcal{N}, t \in \mathcal{T} \end{aligned} \tag{5}$$

$$p_{ij,t}^{RT}, p_{ij,t}^{PB}, v_{ij,t}^{RT}, x_{ij,t} \geqslant 0, \quad \forall i, j \in \mathcal{N}, i \neq j, t \in \mathcal{T} \tag{6}$$

where $c_{ij}$ denotes the cost of relocating a vehicle from zone $i$ to zone $j$. Constraints (2) are capacity constraints that ensure the total number of vehicles dispatched to serve demand or to be relocated does not exceed the available vehicles in each zone during a given time slot. Constraints (3) set the predicted real-time demand, $n_{ij,t}^{RT}(p_{ij,t}^{RT})$, as the upper bound for realized demand, while Constraints (4) guarantee that all predicted pre-booked orders will be fulfilled. Constraints (5) are flow conservation constraints that update vehicle availability for the next time slot $t + 1$, where $\delta_{ji}$ is the travel time from zone $j$ to zone $i$. Constraints (6) define the feasible domain of the decision variables.

After optimizing the reservation prices $\boldsymbol{p}^{PB*}$ in the offline phase, passengers decide whether to pre-book a ride, providing the platform with accurate pre-booked demand $\boldsymbol{d}^{PB}$. However, real-time demand remains uncertain, and the platform lacks knowledge of future real-time demand distributions. Thus, real-time pricing and vehicle relocation decisions must be dynamically adjusted based on demand and supply.

We model the dynamic pricing and vehicle relocation problem as a Markov Decision Process (MDP) and solve it using Deep Reinforcement Learning (DRL). The MDP is defined as a tuple $(\mathcal{S}, \mathcal{A}, \mathcal{R}, \mathcal{P})$, where:

- **State Space** $\mathcal{S}$: Includes pre-booked demand $\boldsymbol{d}_t^{PB}$, real-time demand $\boldsymbol{D}_{t-1}^{RT}$, and vehicle availability $\boldsymbol{q}_t$. A state at time step $t$ is defined as $s_t = \{\boldsymbol{d}_t^{PB}, \boldsymbol{D}_{t-1}^{RT}, \boldsymbol{q}_t\}$.

- **Action Space** $\mathcal{A}$: Consists of pricing $\boldsymbol{p}_t^{RT}$, vehicle dispatching $\boldsymbol{v}_t^{RT}$, and relocation decisions $\boldsymbol{x}_t$. A action at time step $t$ is defined as $a_t = \{\boldsymbol{p}_t^{RT}, \boldsymbol{v}_t^{RT}, \boldsymbol{x}_t\}$.

- **State Transition** $\mathcal{P}$: The transition probability describes how demand evolves based on the actions taken, while vehicle availability changes deterministically according to the relocation and dispatching decisions.

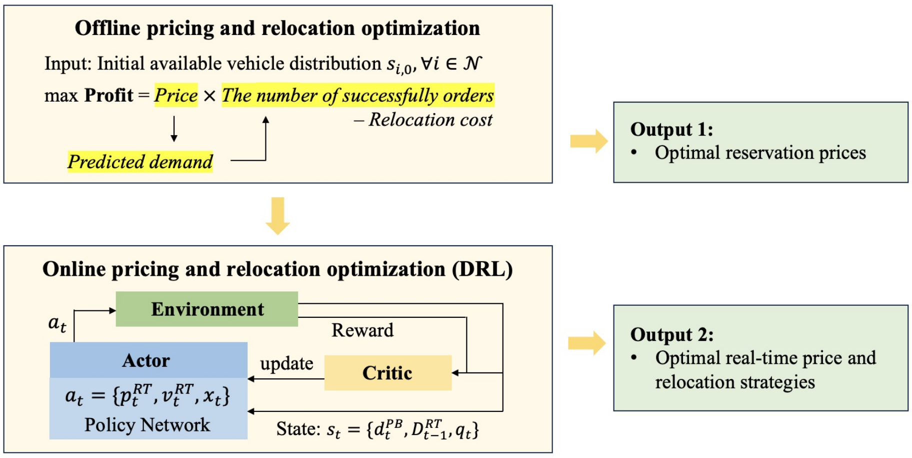


Figure 1 – *An offline-online combined joint pricing and vehicle relocation optimization framework*

- **Reward Function** $\mathcal{R}$: The reward evaluates the platform's profit based on the real-time demand fulfilled, considering revenue from both pre-booked and real-time orders, minus the vehicle relocation costs.

# 3 RESULTS

We conduct a numerical experiment using the historical order dataset provided by Didi Chuxing in Haikou, China, from 1/5/2017 to 31/10/2017 (six months). The study area is divided into approximately 111 hexagonal zones, each with a side length of 1.22 km and an area of 5.16 km$^2$, as shown in Figure 2. This division results in 2,616 OD pairs, with real-time and pre-booked orders counted for each pair every 10 minutes. In total, the dataset includes 87,150 real-time orders and 1,353 pre-booked orders. For the real-time orders, the zones with the top 300 highest order counts were selected, as shown in Figure 3. This selection focuses on regions with higher demand, enabling the observation of significant demand fluctuations, which in turn improves the accuracy of demand forecasting.

Relocation costs are computed based on the Manhattan distance between zones, with a unit cost of 0.5 CNY per kilometer per vehicle. The initial vehicle supply across zones is randomly generated to reflect real-world variability. All experiments were executed on a laptop with a 2.0 GHz quad-core Core i5 CPU.

## 3.1 Comparison of different online optimization

We employ an offline-online optimization framework, where the offline phase predicts the demand for both pre-booked and real-time orders in order to determine the optimal pre-booked prices for the online phase. In the offline phase, linear regression models the demand as a function of price, expressed as $p_{ij,t}^{RT} = \alpha + \beta d_{ij,t}^{RT}$, with parameters $\alpha$ and $\beta$ derived from historical data. This is further enhanced by an LSTM model that captures temporal dependencies, resulting in a hybrid model that improves demand forecasting accuracy. Based on these demand predictions, the offline model sets the pre-booked prices and allocates vehicles to meet the forecasted demand, aiming to maximize profits and minimize unmet orders.

In the online phase, we implement a DRL framework that dynamically adjusts pricing and vehicle relocation decisions based on real-time observations at each time step $t$. The DRL model optimizes the system's performance by considering the current state, including pre-booked

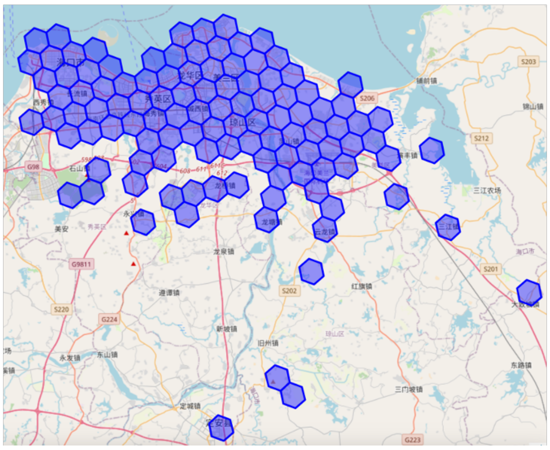

Figure 2 – *Distribution and regional division of orders*

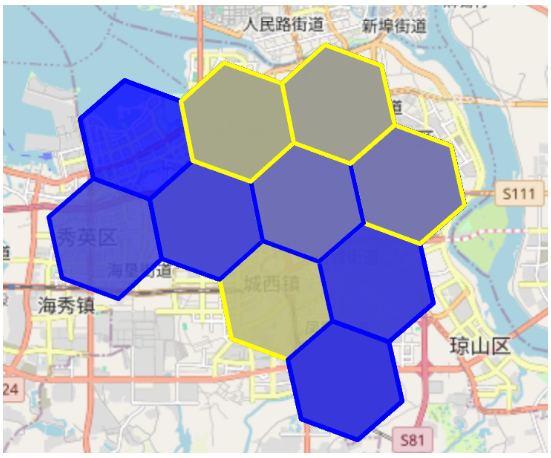


Figure 3 – *Selected zones for numerical experiments*

demand, real-time demand from the previous time step, and vehicle availability. This approach allows the system to adapt to fluctuations in both supply and demand, ensuring operational efficiency.

Results presented in Table 1 show that our DRL-based approach consistently enhances profits across various scenarios, outperforming traditional methods. The training set consists of data from 4/30/2017 to 5/15/2017, while the test set is based on data from 5/16/2017 to 5/21/2017.

Table 1 – *Profit Improvement on Different Operating Hours*

| | Relocation-only($) | Online pricing($) | Improvement(%) |
|---|---|---|---|
| Off-Peak | 67268 | 68723 | 2.16% |
| Peak | 28275 | 29159 | 3.13% |

## 4 DISCUSSION

In this study, we propose a "predict-then-optimize" data-driven framework to optimize pricing and vehicle relocation for ridesourcing platforms, addressing the challenge of managing both pre-booked and real-time orders. To manage demand uncertainty, we develop an offline-online hybrid model where the optimal trip fares for pre-booked orders are determined in the offline phase, while dynamic vehicle relocation and pricing for real-time orders are determined in the online phase using a DRL-based approach. Numerical experiments with real data validate the effectiveness and applicability of the proposed framework.

## References

Engelhardt, Roman, Dandl, Florian, & Bogenberger, Klaus. 2022. Simulating ride-pooling services with pre-booking and on-demand customers. *arXiv preprint arXiv:2210.06972*.

Huang, Jianbin, Huang, Longji, Liu, Meijuan, Li, He, Tan, Qinglin, Ma, Xiaoke, Cui, Jiangtao, & Huang, De-Shuang. 2022. Deep reinforcement learning-based trajectory pricing on ride-hailing platforms. *ACM Transactions on Intelligent Systems and Technology (TIST)*, **13**(3), 1–19.

Yahia, Cesar N, de Veciana, Gustavo, Boyles, Stephen D, Abou Rahal, Jean, & Stecklein, Michael. 2021. Book-ahead & supply management for ridesourcing platforms. *Transportation Research Part C: Emerging Technologies*, **130**, 103266.