Addressing supply and demand heterogeneity: Trip order menus in Ride-hailing platforms

Yue Yang[a], Mohsen Ramezani[a,*]

[a] The University of Sydney, School of Civil Engineering, Sydney, Australia
yyan2437@uni.sydney.edu.au
mohsen.ramezani@sydney.edu.au

Extended abstract submitted for presentation at the 11th Triennial Symposium on Transportation Analysis conference (TRISTAN XI)
June 19-25, 2022, Mauritius Island

December 15, 2021

Keywords: Ride-hailing, Heterogeneity, Matching elasticity

1 INTRODUCTION

In 2011, Uber first launched its ride-hailing service in San Francisco. The rise of Uber paved the way for more and more platforms (i.e. Didi Chuxing and Lyft) to provide ride-hailing services all around the world, reshaping transportation systems in both urban and suburban areas. The essential role of these platforms is 'service providers' to offer value-added services for both demand (i.e., travelers) and supply (i.e., drivers). Ride-hailing platforms are a typical two-sided market, so making matching proposals (Nourinejad & Ramezani, 2016) and setting pricing strategies (Nourinejad & Ramezani, 2020) are two critical components to connect demand and supply sides. In the context of order matching, much work (Zhan *et al.*, 2016, Xu *et al.*, 2018) has been conducted recently assuming that all the drivers working in platforms are full-time employees and service (accept) the dispatched orders unconditionally. After the matching process, the platform charges a fare to travelers and pays a wage to drivers by withholding the commission, which is normally between 15% and 30%, depending on the time, region, and company.

On the supply side in ride-sourcing platforms, drivers make working decisions (i.e. when and where to service) and have heterogeneous market-behavioural patterns (i.e. being full-time or part-time) based on their income, waiting time, etc. Also, travel purpose, the value of time, spatial location, and other factors result in the heterogeneity of the trip demand. Taking into account the heterogeneity in the matching process, the platform can not ensure that a driver will accept a dispatched trip order. To mitigate the uncertainty in one-to-one matching, we consider a relaxed matching pattern that offering each driver a menu of trip orders to choose from. A key issue then is to determine which orders should be listed in the menu of each driver. Such menus are needed to be designed carefully to nudge the self-optimizing drivers' behaviours toward a desirable collective outcome.

In this paper, the above problem is solved by combining modeling driver's choice and designing order menus: i) we investigate the drivers' choice behaviour and model the probability of choosing an order or ignoring the whole order menu. ii) We model the problem as many-to-many matching and develop an algorithm to design order menus. Through numerical experiments on the simulator based on the Manhattan road network, the proposed method could improve the platforms' efficiency and enhance drivers' and travelers' experience.

2 Methodology

2.1 Modeling driver's choice

Let $D = \{d_1, d_2, ..., d_n\}$ denote a set of drivers and $O = \{o_1, o_2, ..., o_m\}$ be a set of trip orders. For an arbitrary pair with a driver $d \in D$ and an order $o \in O$, a utility function is defined as $u_{d,o}$.

$$u_{d,o} = \beta_{0,d} + \beta_{1,d} \cdot f_o - \beta_{2,d} \cdot \tau(l_o^{\text{org}}, l_d) + \beta_{3,d} \cdot V(l_o^{\text{dest}}) \tag{1}$$

where f_o represents o's fare (priced based on the travel distance and travel time), $\tau(l_o^{\text{org}}, l_d)$ indicates the pick-up distance from driver's location, l_d, to order o's origin, l_o^{org}. $V(\cdot)$ is the spatial-value function derived from demand statistics; higher $V(\cdot)$ indicates the place have a higher probability of being matched to the subsequent passenger. There are three main factors considered in the function: (i) order fare, an order with a higher fare will naturally attract the driver (i.e. less probability of decline by the driver). (ii) Pickup distance, a longer pickup distance undermines the driver's willingness to choose an order. (iii) Order destination, an order whose destination is a passenger-hotspot region will be more probable to be selected.

Given an order menu O_d and 'Decline' option c for driver d, the probability of option $o, \forall o \in O_d \cup \{c\}$ is chosen by driver d:

$$p_{d,o} = \frac{e^{u_{d,o}}}{e^{u_{d,c}} + \sum_{o' \in O_d} e^{u_{d,o'}}} \tag{2}$$

where $u_{d,c}$ is a constant (i.e. the expected profit for serving an order) to represent the reserved expectation of driver d by declining the current order menu O_d and wait for the next instance of matching.

2.2 Designing menus

Suppose unserved orders O and idle drivers D are collected at each matching instance. Let $x_{d,o}$ be a binary decision variable that equals 1 if order $o, \forall o \in O$ is listed in the order menus of driver $d, \forall d \in D$, and 0 if not. The probability $\mathbb{P}_o$ of order o selected by at least one driver is:

$$\mathbb{P}_o = 1 - \prod_{d \in D} [1 - p_{d,o}(\{x_{d,o'} | \forall o' \in O\}) \cdot x_{d,o}] \tag{3}$$

where $p_{d,o}(\{x_{d,o'} | \forall o' \in O\})$ (abbreviated as $p_{d,o}$ in the sequel) is the choosing probability that d chooses o if the menu of d is known as $\{x_{d,o'} | \forall o' \in O\}$, and $\prod_{d \in D}[1 - p_{d,o}(\{x_{d,o'} | \forall o' \in O\}) \cdot x_{d,o}]$ represents the probability that none of the drivers choose order o.

In this paper, the objective is maximizing the sum of $\mathbb{P}_o$ over all orders in O so as to maximize the number of orders being responded:

$$\max \sum_{o \in O} \mathbb{P}_o = \max \sum_{o \in O} [1 - \prod_{d \in D} (1 - p_{d,o} x_{d,o})] \tag{4}$$

To achieve the above goal, we abstract idle drivers $d, \forall d \in D$ and unserved orders $o, \forall o \in O$ as two sets of vertices, and all valid pairs (if $x_{d,o} = 1, \forall o \in O, \forall d \in D$) as the set of edges with a weight $u_{d,o}$. After that, the optimization problem can be transformed into determining the optimal edge configuration in the bipartite graph. Given a fully-connected graph $X = \{x_{d,o} = 1 | \forall o \in O, \forall d \in D\}$, we introduce an iterative edge cutting algorithms to achieve the best edge configuration in the bipartite graph.

Once we cut off the edge between o and d and remove o from driver d's orders menus, the choosing probabilities of choosing remaining existing orders o', $o' \in O_d$ are updated. We can define $\delta_{d,o'}$ to capture the change of $p_{d,o'}$:

$$\delta_{d,o'} = p^{\text{update}}_{d,o'} - p_{d,o'}; \forall o' \in O_d/\{o\} \tag{5}$$

where $p^{\text{update}}_{d,o'}$ is a probability that driver d chooses o' in the new graph ($x_{d,o} = 0$).

Furthermore, $\delta_{d,o'}$ leads to a change of $\mathbb{P}_{o'}$:

$$\mathbb{P}^{\text{update}}_{o'} - \mathbb{P}_{o'} = \delta_{d,o'} \cdot \prod_{d' \in D/\{d\}} (1 - p_{d',o'} x_{d',o'}); \forall o' \in O_d/\{o\} \tag{6}$$

where $\mathbb{P}^{\text{update}}_{o'}$ indicates the updated value of $\mathbb{P}_{o'}$ in the new graph ($x_{d,o} = 0$).

On the other hand, $\mathbb{P}_o$ also changes after cutting off the edge between o and d:

$$\mathbb{P}^{\text{update}}_{o} - \mathbb{P}_{o} = -p_{d,o} \cdot \prod_{d' \in D/\{d\}} (1 - p_{d',o} x_{d',o}). \tag{7}$$

Therefore, we can obtain the overall gain Δ of the objective function:

$$\Delta = \sum_{o' \in O_d/\{o\}} (\mathbb{P}^{\text{update}}_{o'} - \mathbb{P}_{o'}) + \mathbb{P}^{\text{update}}_{o} - \mathbb{P}_{o}, \tag{8}$$

positive Δ indicates an improvement of our objective function.

To achieve the optimal edge configuration in the graph, the most straightforward way is to compute Δ for all the edges in the graph, and then cut off the 'best' edge (let $x_{d^*,o^*} = 0$) with the maximal gain $\Delta^{\max}$ for the next iteration. The above process will be repeated until there is no edge with a positive Δ. We call the method Edge Cutting algorithm (see Algorithm 1).

Algorithm 1 Edge Cutting

Input: Unserved orders O, idle drivers D
Output: Order menus X^*
  Initialize a graph $X = \{x_{d,o} = 1 | \forall o \in O, \forall d \in D\}$
  Find out the 'best' edge x_{d^*,o^*} with the maximal Δ^{max}
  while $\Delta^{\max} > 0$ **do**
    Cut off the edge between o^* and d^* (let $x_{d^*,o^*} = 0$)
    Update $p_{d,o}, \forall o \in O, \forall d \in D$ in the graph
    Find out the 'best' edge x_{d^*,o^*} with the maximal Δ^{max}
  end while
  $X^* \leftarrow X$
  return X^*

3 Experiment

In this section, we examine the performance of the proposed method. All the experiments are conducted in a simulation environment in which Manhattan island is considered. The demand data used are Manhattan taxi datasets in December 2020 collected from Yellow Cab's website[1]. The trip order data include order location (origin and destination) and order request time. Each passenger is assigned a matching patience time stochastically drawn from a truncated Gaussian distribution in the range of 0.5 [min] to 1 [min] with a mean of 0.75 [min] and standard deviation of 0.15 [min]. An order will be canceled if not being matched or picked up within the matching patience time. 300 drivers are initially generated in the road network at 07:00 AM randomly. The number and location of new arriving drivers are considered to be stochastic and time-varying.

[1]https://www1.nyc.gov/site/tlc/about/tlc-trip-record-data.page

Furthermore, drivers are assumed to be impatient; they will leave the network once they receive no matching for 30 [min]. On the other hand, the platform kick out the drivers if they take more than 10 [min] to choose a ride.

Every ten seconds, the platform provides an order menu to each idle driver. Drivers choose a ride among the offered order menu (with the option decline) based on the utility and logit model (Eq. 1 and 2). Subsequently, the platform collects the drivers' decisions and assigns the nearest responded driver to each order. Note that Idle drivers are parked at the last drop-off location.

The performance of the following benchmark methods are evaluated: i) **One-to-one**: This method is commonly employed in 'Order Dispatching' or 'Ride Matching' scenarios. ii) **Global**: All orders are listed in the menus of idle drivers. iii) **Local**: An order is listed in the menu of drivers whose pick-up distance is less than 2 kilometers.

Table 1 – *The matching results under different strategies to design order menus.*

Method	Avg. Response	Avg. Cancellation	Avg. Response time [s]	Avg. Occupied rate
One-to-one	4679.7 (63.6%)	2640.1 (35.9%)	13.1	47.2%
Global	3467.3 (47.1%)	3840.4 (52.2%)	17.7	43.4%
Local	5327.8 (72.5%)	1998.5 (27.2%)	12.4	51.3%
Edge Cutting	5767.3 (78.5%)	1563.4 (21.2%)	11.5	53.7%

Table 1 summarizes the outcomes of different methods implemented in the simulator. The Global method leads to a worse platform efficiency due to the concentration of the drivers' willingness on several attractive orders (over-competition on a few high-utility orders). Contrarily, the Local and Edge Cutting methods can achieve more promising results. It is worth noticing that though the Edge Cutting algorithm is designed for optimizing average response in our problem, it achieves the best performance on all evaluation metrics with 78.5% average response rate, 21.2% average cancellation rate, 11.5 [s] average response time and 53.7% average occupied rate of the vehicles.

4 Conclusion

Considering the heterogeneity of both supply and demand in the matching process, we propose a method to design the trip order menus in peer-to-peer ride-hailing platforms. Through extensive experiments, the proposed method enhance both passengers' and drivers' experience and improve platform efficiency

References

Nourinejad, Mehdi, & Ramezani, Mohsen. 2016. Developing a large-scale taxi dispatching system for urban networks. *Pages 441–446 of: 2016 IEEE 19th International Conference on Intelligent Transportation Systems (ITSC)*. IEEE.

Nourinejad, Mehdi, & Ramezani, Mohsen. 2020. Ride-Sourcing modeling and pricing in non-equilibrium two-sided markets. *Transportation Research Part B: Methodological*, **132**, 340–357.

Xu, Zhe, Li, Zhixin, Guan, Qingwen, Zhang, Dingshui, Li, Qiang, Nan, Junxiao, Liu, Chunyang, Bian, Wei, & Ye, Jieping. 2018. Large-scale order dispatch in on-demand ride-hailing platforms: A learning and planning approach. *Pages 905–913 of: Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*.

Zhan, Xianyuan, Qian, Xinwu, & Ukkusuri, Satish V. 2016. A graph-based approach to measuring the efficiency of an urban taxi service system. *IEEE Transactions on Intelligent Transportation Systems*, **17**(9), 2479–2489.