6-2019

Matching passengers and drivers with multiple objectives in ride sharing markets

Guodong LYU

Chung Piaw TEO

Wangchi CHEUNG

Hai WANG
Singapore Management University, haiwang@smu.edu.sg

Citation

LYU, Guodong; TEO, Chung Piaw; CHEUNG, Wangchi; and WANG, Hai. Matching passengers and drivers with multiple objectives in ride sharing markets. (2019). *Proceedings of the 10th Triennial Symposium on Transportation Analysis, Hamilton Island, Australia, 2019 June 17-21*. 1-4.
Available at: https://ink.library.smu.edu.sg/sis_research/6873

Matching Passengers and Drivers with Multiple Objectives in Ride Sharing Markets

Guodong Lyu, Chung Piaw Teo

Department of Analytics and Operations

National University of Singapore

Wangchi Cheung

Department of Industrial Systems Engineering and Management

National University of Singapore

Hai Wang

School of Information Systems

Singapore Management University

Email: haiwang@smu.edu.sg

In many cities in the world, ride sharing companies, such as Uber, Didi, Grab and Lyft, have been able to leverage on Internet-based platforms to conduct online decision making to connect passengers and drivers. These online platforms facilitate the integration of passengers and drivers' mobility data on smart phones in real-time, which enables a convenient matching between demand and supply in real time. These clear operational advantages have motivated many similar shared service business models in the public transportation arena, and have been a disruptive force to the traditional taxi industry.

Matching passengers (demand) with drivers (supply) in real-time is a challenging problem for the ride sharing platforms. Greedy policy, as a common used benchmark matching policy, assign passenger to the "nearest" available driver, based on the pick-up time estimated from each driver's location and surrounding traffic conditions. Note that the pick-up time of the assigned driver affects whether a passenger will cancel the booking, or show up at the pick up location. In reality, platforms also consider many other objectives in the matching policy. One important objective is the rating for drivers. Uber, for instance, uses rating provided by passengers to rate the drivers, and booted those drivers whose rating fall below a threshold from their system. Platforms usually give priority to drivers with higher ratings in the matching policy, especially during the off-peak hours with sufficient supply. Such preference to drivers with higher rating could encourage drivers to provide better service, which will also improve the overal service quality in the system. Another important objective is the passenger revenue (i.e., the order estiamted fare). Platforms also give priority to passengers with higher revenue in the matching policy, especially during the peak hours with a large number of passengers. Such preference to passengers with higher revenue could bring more revenue and profit to the platforms.

In general, the matching decisions between drivers and passengers are supposed to take the trade-offs between multiple objectives into account. Although piles of efforts have been devoted to designing matching policies for the two-sided sharing market, the majority of these works focused on a single-objective optimization problem. For example, Zhang et al. (2017) [1] develop a batch matching system, with the objective to maximize the driver acceptance rate for each order. Different from the traditional one-order-to-one-driver matching mechanism, they dispatch each order to multiple drivers and let drivers compete for the order. Hu and Zhou (2016) [2] study the dynamic matching control of a two-sided, discrete-time matching system in which both the supply and demand may leave the platform if the wait time before getting passengers or drivers are too long, with the objective to maximize the expected total discounted profit. Ozkan and Ward (2016) [3] propose a linear programming based matching policy that accounts for temporal changing demand and supply and customer patience, with the objective of maximizing the overall number of passengers being served. Wang et al. (2017) [4] introduce the concept of stability in dynamic ride sharing and provide mathematical programming approaches to solve stable and nearly stable ride-share matching problems, with the objective of minimizing the pick-up detour distance. However, few studies shed light on the multi-objective matching policy in the ride sharing markets.

In this paper, we study the matching problem for ride-sharing platform with multiple objectives, and design an online matching policy that simultaneously achieves multiple objectives in a balanced manner. More precisely, we aim to achieve a solution that has the smallest deviation, based on some pre-determined distance function, to an "utopia point", i.e., an ideal solution maximizing the performance of all objectives, but is otherwise non-attainable at the same time. The obtained solution with shortest deviation to the target is called the "compromise solution". We apply the online policy in ride sharing market settings, and provide an online matching policy that simultaneously incorporates driver service scores (driver ratings), pick-up distances and passenger revenues. To be more specific, the platforms would want to dispatch more passenger orders to drivers with higher service rating. This helps to retain the better drivers in the system, and provide better service experience to the customers. However, this could not come without sacrificing the average pick-up distance between dispatched drivers and passengers. Moreover, the platform needs to manage the impact on the bottom line - longer waits lead to lower answer count (passengers drop the bookings) and lower revenue. To balance these different Key Performance Indexes (KPIs), three key considerations need to be taken into account to design the matching policies in these markets: (1) Passengers with higher revenues should be served with higher priority; (2) Passengers' waiting time for pick-up should be as small as possible; (3) Drivers with higher scores should be dispatched with higher priority.

Note that the traditional approach to multi-objective optimization problem entails a delicate selection of weighting function to aggregate the multiple objectives into a single one, and the central issue there is the choice of the weighting function to be used for aggregation. Our approach exploits the multiple period setting, and the existence of natural performance targets (i.e., the utopia point), to develop an adaptive weighting function that learns from historical performance to drive the algorithm towards the compromise solution. Our detailed numerical studies on the driver dispatching problem show that this approach is able to learn from data the appropriate weighting function that can be used in each period to guide the system towards a good matching solution.

We extracted real world data from Didi Chuxing, the largest on-demand ride sharing platform in China, and conducted industrial implementation of the proposed matching policy. Compared to legacy policies currently in use, such as the weighted average policy (Legacy Policy) or the "closest distance" policy (CD Policy), we observe that all parties in the ride-sharing eco-system, from drivers, passengers, to the platform, are better off under our proposed online matching policy generating the compromise solution (CM Policy): (1) drivers with higher service scores are dispatched with more orders; (2) passengers are more likely to be matched to drivers with higher service scores, and passengers with higher revenues (longer travel distances) are served with higher answer rates; (3) the platform obtains a higher revenue and better long-term brand reputation. For instance, we observe that more jobs are assigned to drivers with higher service quality. Figure 1 demonstrates that expected total revenue earned by drivers with higher service scores (e.g., higher than 101) increases under the CM policy. We also find that the revenue increment for these drivers is indeed due to more orders being dispatched to them. This outcome would motivate drivers to increase their service score by providing better ride sharing service to passengers. We observe a decreasing trend in total revenue for these drivers with extreme high service scores. One possible explanation is that a large proportion of drivers are part-time and their revenue also depends on their total business hours (i.e., active time as a driver on the platform). The dataset reveals this pattern: these drivers with service scores in the interval [98, 108] are more active than the ones with scores in the interval [109,116]. Even so, our proposed policy dispatches more orders to these drivers with higher service scores consistently. As a side effect, the total revenue obtained by the platform during the whole day under the our policy also increases by 0.26% and 0.56% in two cities, respectively.

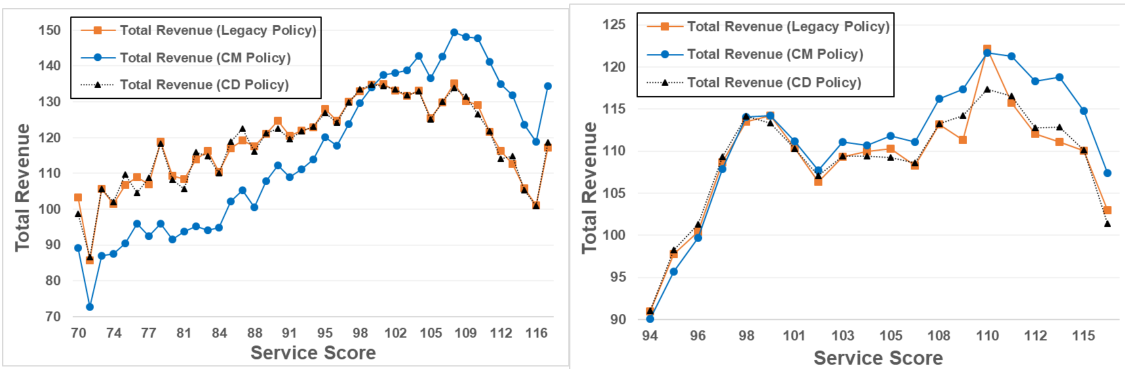

(a) City A

(b) City B

Figure 1: Driver Service Score vs. Driver Revenue

References

[1] Zhang, Lingyu and Hu, Tao and Min, Yue and Wu, Guobin and Zhang, Junying and Feng, Pengcheng and Gong, Pinghua and Ye, Jieping, "A taxi order dispatch model based on combinatorial optimization", *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2151–2159 (2017).

[2] Hu, Ming and Zhou, Yun, "Dynamic type matching", *Working Paper*, (2016).

[3] Ozkan, Erhun and Ward, Amy, "Dynamic matching for real-time ridesharing", *Working Paper*, (2017).

[4] Wang, Xing and Agatz, Niels and Erera, Alan, "Stable matching for dynamic ride-sharing systems", *Transportation Science*, forthcoming (2017).