Investigating Methods of Balancing Inequality and Efficiency in Ride Pooling

Naveen Raman
University of Maryland, College Park
nraman1@umd.edu

Abstract

Our research focuses on developing matching policies that match drivers and riders for ride-pooling services. We aim to develop policies that balance efficiency and various forms of fairness. We did this through two methods: new matching algorithms that include a fairness term in the objective function, and income redistribution methods based on the Shapley value of a driver. I tested these methods on New York City Taxicab data to evaluate their performance and found that they succeed in reducing certain forms of fairness.

Introduction

Ride pooling services have exploded in popularity in recent years, exemplified by the growth of Uber Pool, Lyft Line, and Didi Chuxing. Ride pooling services match riders and drivers according to a policy that maximizes some objective function, such as profit. The policy determines how different rider-driver combinations are weighted, in effect, ranking which drivers to match with which riders. However, these policies have led to fairness issues, such as a gender gap (Cook et al. 2018) and differing pickup rates based on race (Brown 2018).

We deal with the problem of fairness in ride-*pooling* services, which has not been dealt with prior, is more difficult than fairness in rideshare due to the increase in combinations of riders and drivers. We pursue two methods of incorporating fairness: through new matching policies and income redistribution methods. We make the following contributions:

1. Investigation of how different parameters, such as the number of riders and drivers, affect fairness and profitability
2. New policies that aim to balance fairness with profitability. In particular, we propose two new policies, one of which aims to minimize a form of inequality on the rider side, and another that aims to minimize on the driver side
3. An income redistribution strategy that reduces income inequality while avoiding the free-rider problem.

Related Work

We build on prior work in matching in rideshare and ride-pooling, and fairness in rideshare. Past work on matching helps us develop a framework to evaluate policies.

Matching in Rideshare

Prior work has viewed the rideshare and ride-pool matching problems as an instance of the Markov Decision Process framework. Much of this work has approached the problem from an online matching perspective, where deep learning is used offline to learn the value of the state, and is used to optimize matches online (Lin et al. 2018). To deal with the complexity of the ride-pool matching problem, past work has used Approximate Dynamic Programming has been used in conjunction with deep learning to efficiently match riders and drivers (Shah, Lowalekar, and Varakantham 2020).

Fairness

Fairness in rideshare and ride-pooling has been studied from both an algorithmic perspective and a social science perspective. Past research discusses a trade-off between profit and certain definitions of fairness in rideshare, which can be adjusted through a parameter (Nanda et al. 2020). There have been reports of disparate treatment by rideshare companies for certain sub-populations in terms of differences in wait-time between black and non-black riders (Brown 2018).

My Contributions

My main contributions to the project were developing the policies, implementing the policies, and writing up the results. I extended a framework to test policies (Shah, Lowalekar, and Varakantham 2020) and developed new policies that balanced fairness and efficiency. To determine what policies to use, I talked with my mentor, and read background work to see what types of objective functions might work best. Afterwards, I ran experiments with these policies, implemented income redistribution, and then plotted the data for both policies and redistribution.

Policy Experiments

We conducted two types of experiments, the first testing different policies, and the second testing different income redistribution methods. To evaluate a policy, I extended

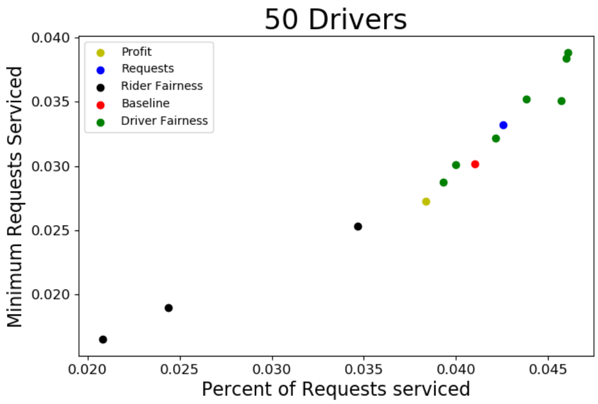

Figure 1: We compare policies at 50 drivers, with each dot representing a combination of objective function and hyper-parameter. We find that with 50 total drivers, objective functions that minimize the spread of income also achieve the most rider requests both overall, and also in the worst-off neighborhood. The methods optimizing for inequality outperform state of the art methods, which optimize for the number of requests, are titled "Requests" on the chart.

the Approximate Dynamic Programming framework (Shah, Lowalekar, and Varakantham 2020), and used deep learning to learn the value of a state for each policy.

We compare policies on the rider and driver side, and find that when the number of riders is much more than the number of drivers, policies that aim to minimize the variance of wages, also end up serving the most requests, both overall, and in the worst-off neighborhoods (Figure 1). This means that certain definitions of inequality are aligned with maximizing the number of riders serviced, and so rider and driver-side inequality could be accomplished without sacrificing profitability. These policies can then be applied to real-world ride-pooling systems to measure their effects upon inequality and profit in a real setting.

Income Redistribution Experiments

We propose a second method to encode a variant of driver-side fairness, though instead of encoding it at the matching level, we instead attempt to redistribute income at the end of the day. Each driver takes a certain commission from each rider, and the rest is redistributed at the end of each day.

We introduce two methods of determining the value of a driver, the first being the amount of time a driver spends driving, and the second being the driver's Shapley Value. The Shapley value of driver is a way of measuring the marginal contribution of a driver to every subset of drivers, which helps assess the true contributions of each driver, independent of the presence of other drivers. Using either of these methods of value, the income forgone by each driver is redistributed proportionally to how much less than their value each driver earned. For example, if there's a total of \$9 forgone by all drivers, and one driver earns \$10 less than their value, while the other earns \$5 less than their value, then the first driver will receive \$6 and the second driver will receive \$3.

We vary the risk parameter, the number of drivers, and the policy used to evaluate how income redistribution performs. We evaluate based on the spread of income and the correlation between value and final payment. If the value and final payment are uncorrelated, then we run into the free-rider problem, where all drivers get paid the same amount, regardless of value. We aim to minimize the spread of income while avoiding the free-rider problem.

We find that, for certain risk thresholds, we can keep a high gain, while minimizing income inequality. This works regardless of whether the method used to determine the value of a driver is the time spent driving or the Shapley value.

Next Steps

Next steps for my research would be to investigate other forms of inequality within ride-pooling. In particular, differential treatment based on race (Brown 2018) can partially be attributed to differences in surge pricing (Stark and Diakopoulos 2016). In particular, surge pricing encourages drivers to service white neighborhoods. I would like to investigate whether Uber's surge pricing propagates existing socioeconomic differences between neighborhoods, and if there's a way to modify surge pricing in a way to avoid this. Doing this would first require learning some type of demand function, in effect, to determine at what prices riders would reject rides. This could be done using a supervised learning algorithm based on real-world data using data on rider preferences. From there, I would research how Uber's surge pricing algorithm works, implement it, and see how it interacts with some of the policies explored in this paper. I would then change some details of the algorithm to see whether we can avoid discrimination while preserving profit.

References

Brown, A. E. 2018. *Ridehail revolution: Ridehail travel and equity in Los Angeles*. Ph.D. thesis, UCLA.

Cook, C.; Diamond, R.; Hall, J.; List, J. A.; and Oyer, P. 2018. The gender earnings gap in the gig economy: Evidence from over a million rideshare drivers. Technical report, National Bureau of Economic Research.

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. ACM.

Nanda, Anant Xu, P.; Sankararaman, K. A.; Dickerson, J. P.; and Srinivasan, A. 2020. Balancing the Tradeoff between Profit and Fairness in Rideshare Platforms During High-Demand Hours. In *Conference on Artificial Intelligence*.

Shah, S.; Lowalekar, M.; and Varakantham, P. 2020. Neural Approximate Dynamic Programming for On-Demand Ride-Pooling. In *Conference on Artificial Intelligence (AAAI)*. Full version: arXiv:1911.08842.

Stark, J.; and Diakopoulos, N. 2016. Uber seems to offer better service in areas with more white people. That raises some tough questions. *The Washington Post* .