
Replication study of "Data-Driven Methods for Balancing Fairness and Efficiency in Ride-Pooling"

Anonymous Author(s)

Affiliation

Address

email

Reproducibility Summary

1

2 **Scope of Reproducibility**

3 We evaluate the following claims related to fairness-based objective functions presented in [4] : (1) For the four objective
4 functions, the success rate in the worst-off neighborhood increases monotonically with respect to the overall success rate.
5 (2) The proposed objective functions do not lead to a higher income for the lowest-earning drivers, nor a higher total
6 income, compared to a request-maximizing objective function. (3) The driver-side fairness objective can outperform a
7 request-maximizing objective in terms of overall success rate and success rate in the worst-off neighborhood. This
8 means that this objective, whilst reducing the spread of income, also positively impacts rider fairness and profitability.

9 **Methodology**

10 The code provided by [4] was used as a base for our re-implementation in PyTorch. We evaluate the claims by
11 the original authors by (a) replicating their experiments, (b) testing for sensitivity to a different value estimator, (c)
12 examining sensitivity to changes in the preprocessing method, and (d) testing for generalizability by applying their
13 method to a different dataset.

14 **Results**

15 We reproduced the first claim since we observed the same monotonic increase of the success rate in the worst-off
16 neighborhood with respect to the overall success rate. The second claim we did not reproduce, since we found that
17 the driver-side fairness objective function obtains a higher income for the lowest-earning drivers than the request-
18 maximizing objective function. We reproduced the third claim, since the driver-side objective function performs best in
19 terms of overall success rate and success rate in the worst-off neighborhood, and also reduces the spread of income.
20 Changes of the value estimator, preprocessing method and even dataset all led to consistent results regarding these
21 claims.

22 **What was easy**

23 The paper is written engagingly and the theoretical sections, in particular, give a clear description of the problem setup
24 and objectives. The paper is also accompanied by an open-source code base, which supports reproduction efforts.

25 **What was difficult**

26 The provided code lacks a script to preprocess raw data, which is required to reproduce the experiment, nor was
27 the preprocessed data openly available. Additionally, complex code structure and scarce commenting complicated
28 replication.

29 **Communication with original authors**

30 Due to the absence of preprocessed data, we contacted the authors, who quickly provided the requested data.

31 1 Introduction

32 Ride-pooling platforms match independent drivers with multiple riders. This matching is performed by machine
 33 learning algorithms, which are designed to maximize company profit. The profit motive behind these algorithms can
 34 cause unfairness among drivers and riders, for instance by unequally distributing rides between drivers, or by servicing
 35 requests originating in some neighborhoods at lower rates [2, 3, 4].

36 The paper *Data-Driven Methods for Balancing Fairness and Efficiency in Ride-Pooling* [4] (henceforth referred to as
 37 “the paper” or “the authors”) explores the tradeoff between rider fairness, driver fairness, and total income (i.e. company
 38 profitability) by measuring fairness and profitability metrics across simulations using four objective functions. These
 39 objective functions are maximized by an algorithm that matches rider requests to drivers, using actual request data from
 40 the New York Yellow Taxi dataset. The following metrics are studied: a) the percentage of all requests that are serviced
 41 (the *overall success rate*), b) the success rate in the neighborhood with the lowest local success rate (the *success rate in*
 42 *the worst-off neighborhood*), and c) the distribution of income across drivers. The overall success rate is used as a proxy
 43 for company profitability, whereas the success rate in the worst-off neighborhood and the income of the least-earning
 44 drivers are measures of rider and driver fairness, respectively.

45 The matching algorithm, introduced by [6], uses a Markov decision process (MDP) in combination with a neural value
 46 estimator to match rides to drivers non-myopically, that is, with awareness of future events that could impact the value
 47 of a match. The paper’s algorithm requires a strongly connected graph (i.e. street network) on which the taxis operate,
 48 precomputed routes and travel times between all pairs of nodes (i.e. intersections), and a dataset of requests, each
 49 containing an origination node, a destination node, and the time the request was issued.

50 The paper compares two profitability-focused objective functions and two fairness-focused objective functions. The
 51 profitability-focused request objective maximizes the total number of requests serviced during the simulation, given by
 52 the sum of ongoing requests p_i and completed requests s_i , for each driver i :

$$O_{request}(R, W) = \sum_{i=1}^n (|p_i| + |s_i|), \quad (1)$$

53 where R and W are sets respectively containing the states of the drivers, and all previously unaccepted and accepted
 54 requests. The profitability-focused income objective maximizes the total income of all drivers:

$$O_{income}(R, W) = \sum_{i=1}^n \pi_i, \quad (2)$$

55 where π_i is the income of driver i , which is made up of a constant part and a variable part that depends on the distance
 56 of the trip. The rider fairness objective maximizes profit whilst minimizing the variance of the success rate across all
 57 neighborhoods j :

$$O_{rider}(R, W) = -\lambda Var\left(\frac{h_j}{k_j}\right) + \sum_{i=1}^n \pi_i, \quad (3)$$

58 where h_j is the number of serviced requests originating in neighborhood j , k_j is the total number of requests originating
 59 in j , and λ a hyperparameter moderating the regularization. The driver-side fairness objective maximizes profit whilst
 60 minimizing the variance of the income across all drivers i :

$$O_{driver}(R, W) = -\lambda Var(\pi_i) + \sum_{i=1}^n \pi_i. \quad (4)$$

61 2 Scope of reproducibility

62 We focus our reproducibility study on the claims related to the fairness-based objective stated in the previous section.
 63 These claims (henceforth referred to as *the claims*) can be summarized as follows:

- 64 1. For the four objective functions, the success rate in the worst-off neighborhood increases monotonically with
 65 respect to the overall success rate.

- 66 2. The proposed objective functions do not lead to a higher income for the lowest-earning drivers, nor a higher
67 total income, compared to a request-maximizing objective function. This extends claim 1, suggesting that a
68 profit motive generally leads to increased fairness for drivers and riders alike.
- 69 3. The driver-side fairness objective (4) can outperform a request-maximizing objective (1) in terms of overall
70 success rate and success rate in the worst-off neighborhood. This means that this objective while reducing the
71 spread of income, also positively impacts rider fairness and profitability.

72 The authors have demonstrated claim 3 to hold only for 50 drivers. For 200 drivers they find the opposite, namely that
73 equation (1) outperforms equation (4) in terms of fairness as well as profitability metrics.

74 We evaluate these claims by testing whether they still hold under a variety of modifications to the experimental setup.
75 We conduct four experiments that deviate increasingly from the exact setup of the paper: (a) the experiment of the paper
76 is reproduced using the author’s preprocessed data, (b) the non-myopic neural value estimator is replaced by a myopic
77 greedy estimator, (c) the experiment of the paper is replicated using data generated with our own preprocessing method,
78 and (d) the experiment is applied to a different dataset (New York Green Taxi dataset), using our own preprocessing
79 method. Note that all experiments except for (b) use the neural value estimators also used in the original paper.

80 3 Methodology

81 This section explains the methodology used for the four experiments we carried out. An overview of the used code
82 is provided, followed by descriptions of the model, the datasets and the hyperparameters. Finally, we outline the
83 experimental setup and state the computational resources needed to perform the experiments.

84 3.1 Code

85 The code provided by the authors, which is largely based on the code by [6], was used as a base for our re-implementation
86 in PyTorch. The provided code was sufficient to reproduce the experiments performed in the original paper, after the
87 authors emailed the preprocessed data that they used in the paper. This data consists of the graph of Manhattan, travel
88 times, routes, and rider requests from the New York Yellow Taxi dataset mapped to nodes on the graph. Neither the
89 code to generate this data nor the data itself are publicly available.

90 To solve the integer linear problem that determines which set of actions is assigned to which driver, the original code
91 uses the callable library CPLEX 12.8. However, at the time of this replication study, the free edition of CPLEX does
92 not suffice to train the required models, so the no-cost academic edition of CPLEX 20.1¹ was used instead.

93 In addition to an exact reproduction, we examine the paper’s claims’ robustness against a different method of prepro-
94 cessing the same raw data. The paper uses the method described in [5] to preprocess the raw trip data. However, this
95 method is computationally demanding and complex to implement. We developed an algorithm that generates routes and
96 travel times on a graph, but that differs in two ways from [5]. First, we use travel time estimates from OpenStreetMap
97 (OSM), corrected by a multiplication constant, equal to the mean ratio between the actual travel time and the OSM
98 estimate of the travel time, computed over all trips in the dataset.

99 Second, unlike the method used in the paper, we do not compute the Dijkstra algorithm for each pair of nodes, which is
100 an $O(n^2)$ approach, where n is the number of nodes in the graph. For the street network of Manhattan ($n \approx 4000$), this
101 can take days on a typical laptop without GPU acceleration. Instead, our routing algorithm invokes Dijkstra on a total
102 of $\approx 500,000$ pairs of randomly sampled nodes, which yields a coverage of $\approx 60\%$ of all $n^2 = 16\text{M}$ routes, because
103 all subroutes of each route are also optimal routes. The optimal routes between the remaining 40% of node pairs are
104 approximated by setting each remaining route (n, m) equal to the concatenation of subroutes (n, p) and (p, m) , for a
105 predecessor p selected from a set of predecessors of node m , for which routes (n, p) and (p, m) are known and yield
106 the lowest total travel time.

107 Lastly, to test the generalizability of the methods used in the paper, we performed additional simulations using trips in
108 Brooklyn from the New York Green Taxi dataset. The algorithm described above is used to generate routes and travel
109 times between all nodes, given the graph of Brooklyn’s street network.

¹<https://pypi.org/project/cplex/>

110 3.2 Model descriptions

111 The original paper has adapted the model that incorporates an MDP to assign a set of actions to each vehicle from [6].
112 An overview of this algorithm is provided in Appendix A. The model makes use of a neural value estimator that assigns
113 a value to each possible set of actions. In the original paper, the objective functions (equations 1, 2, 3, 4) that this model
114 aims to maximize are varied, along with the number of drivers (50 and 200 drivers are used).

115 The input to the neural value estimator [6] is composed of the current location and path of the vehicle, the permissible
116 delay, the current epoch, the number of other vehicles in the vicinity and the number of requests that were placed in
117 the current epoch. By using an LSTM, this value estimator can take into account non-myopic considerations like the
118 possibility that a future rider request will appear along the route of a current rider request, therefore increasing the
119 value of the current request. The location and path features are embedded using pretrained embeddings, which were
120 computed by a separate network that was trained to predict travel times between any two nodes in the graph.

121 3.3 Datasets

122 The original experiments are conducted on the New York Taxi dataset [1], available at ², consisting of pickup and
123 drop-off locations and times for Yellow Taxi passengers from March 23rd to April 1st and from April 4th to April 8th,
124 2016.

125 The generalization experiments are conducted on the New York Green Taxi dataset ³, consisting of pickup and drop-off
126 locations and corresponding times for Green Taxi passengers in the months February, March and June. We take the
127 subset of trips in Brooklyn, to keep the size of the graph manageable and similar to the graph of Manhattan used in the
128 original paper. Whereas Yellow taxis in practice only serve the business district of downtown Manhattan, the Green taxis
129 serve all of New York. This makes the Green Taxi dataset a good choice to test the generalization of results concerning
130 rider-side metrics, because of differences in affluence between various districts of New York. In addition, Brooklyn is a
131 different geographical area and the Green Taxi dataset contains a different distribution of requests, compared to the
132 Yellow Taxi dataset.

133 3.4 Hyperparameters

134 No hyperparameter search was done for this replication study. Instead, the hyperparameters provided by [4] were used
135 to adhere to the original experimental setup as much as possible. This meant that for the driver-side fairness objective
136 function we set λ to be $\frac{4}{6}$, and for the rider-side fairness objective function, λ was set to be 10^9 . Additionally, we set the
137 constant costs of a ride to \$5, the capacity of a car to 4 riders, the number of neighborhoods to 10, and we batch rider
138 requests per minute.

139 3.5 Experimental setup

140 For each experiment containing a neural value estimator (section 2), the model is trained for each combination of
141 objective function (equations 1, 2, 3, 4) and number of drivers ($\in \{50, 200\}$). For the specific packages and their
142 versions that were used to obtain our results, we refer to our codebase ⁴. Reproducing the experiments by [4] means
143 that we trained all models on 3 days of data, except for the rider-side fairness objective (equation 3) models, which
144 were trained on 2 days of data. All models were evaluated on one day of data.

145 In order to evaluate the claims, the following metrics are of interest: the overall success rate, the success rate in the
146 worst-off neighborhood (also called *minimum request success rate*), and the income per driver. To compare the objective
147 functions, three plots are generated per experiment, in line with the reporting of the original paper: 1) the minimum
148 request success rate as a function of the overall success rate for the simulation with 50 drivers, 2) the same plot for
149 the simulation with 200 drivers, and 3) the distribution of income across drivers for each objective function for the
150 simulation with 200 drivers.

151 During the exact reproduction of the paper’s experiments, we observed that the neural model is not trained (seemingly
152 inadvertently) in the experiment with 50 drivers. The number of training examples for the neural net grows in the

²<https://www1.nyc.gov/site/tlc/about/tlc-trip-record-data.page>

³<https://www1.nyc.gov/site/tlc/about/tlc-trip-record-data.page>

⁴<https://anonymous.4open.science/r/rideshare-replication-815C>

153 number of drivers and the number of days of training data. The authors included a minimum threshold of training
154 examples, which is not met by the experiment with 50 drivers. However, the algorithm silently executes, therefore using
155 an untrained randomly-initialized neural network. Importantly, the value assigned to each action is a linear combination
156 of a deterministic term and the output of the (untrained) neural net. Therefore, even though the network’s outputs are
157 random in the case of 50 drivers, the computed values are not. This finding motivated the experiment where the neural
158 value estimator is replaced by a greedy value estimator.

159 To test the claims under a different data preprocessing method, we run an experiment using the same raw data as the
160 original paper, but with our self-developed preprocessing method (i.e. a graph of Manhattan with routes and travel times
161 computed as outlined in section 3.1). Apart from the data preprocessing, this experiment is identical to the experiment
162 in the paper. With this, we aim to investigate the sensitivity of the claims and the method used by [4] to different data
163 preprocessing methods and in particular to a different travel time computation, which is inherently noisy.

164 Our final experiment uses the Green Taxi dataset and our own preprocessing method in order to study the generalizability
165 of the work by [4]. The models for the Green Taxi dataset were trained on 6 days of data, except for the rider-side
166 fairness objective (equation 3) models, which were trained on 4 days of data. All these models were evaluated on two
167 days of data. For all experiments using our preprocessing method, two new sets of pretrained embeddings were trained
168 for the corresponding graphs of Manhattan and Brooklyn.

169 **3.6 Computational requirements**

170 All experiments were run on one Nvidia GeForce 1080Ti GPU. Training and evaluating a model on the Yellow Taxi
171 dataset for 50 drivers took roughly 2.5 hours, whereas training and evaluating one for 200 drivers took 6-7.5 hours. For
172 the rider-fairness objective function models, which were only trained on two days of data, runtimes were two-thirds of
173 this. The experiments that use a greedy non-neural value estimator took an hour less for 200 drivers and roughly the
174 same amount of time for 50 drivers. For the Green Taxi dataset, more days of data were used to train on, since this
175 dataset contains fewer requests per day. Here, the model took roughly 1 hour to train and evaluate for 50 drivers and
176 10 hours for 200 drivers. Computing embeddings of size 100 for all ≈ 4000 locations took approximately 10 hours.
177 Finding shortest paths and computing travel times between all pairs of locations took 1.5 hours on a laptop with an Intel
178 i5 processor⁵.

179 **4 Results**

180 In this section, we present the results of the reproduction of the original paper, the replacement of the neural value
181 estimator by the greedy non-neural value estimator, the replication using our data preprocessing method, and the
182 generalization experiment using the Green Taxi dataset. For each experiment, we evaluate whether the claims listed in
183 Section 2 are supported by the presented results.

184 **4.1 Reproduction of the original experiment**

185 The first claim states that objective functions which improve the overall success rate for riders also improve the success
186 rate in the worst-off neighborhood. The figures that the authors use to support this are displayed in Appendix B and our
187 reproduced results in Figures 1a and 1b. If objective functions that improve the overall success rate also improve the
188 success rate in the worst-off neighborhood, we should see the success rate in the worst-off neighborhood monotonically
189 increase as the overall success rate increases. This is indeed a pattern that we see in the figures from the original paper,
190 as well as in our figures.

191 Second, the authors claim that no objective function raises wages for the lowest-earning drivers or raises the total income,
192 compared to the requests objective function (equation 1). In the original paper, results related to the income distribution
193 used to support this claim are only obtained for 200 drivers, as can be seen in Appendix B. In our reproduction the
194 function that maximizes total income at 200 drivers is also the requests objective function. However, the maximum
195 income for the least-earning drivers is obtained with the driver-side fairness objective function, as displayed in Figure
196 1c.

⁵Specifically, a MacBook Pro (13-inch, 2020) with 1.4 GHz Quad-Core Intel Core i5 CPU

197 The final claim states that the driver-side fairness objective function (equation 4) outperforms the requests objective
 198 function (equation 1) in terms of overall success rate, success rate in the worst-off neighborhood, and reducing the
 199 spread of income. These first two results are only obtained for 50 drivers, as seen in the figures in Appendix B, and the
 200 last result only for 200 drivers. We reproduce all three of these results. In Figure 1a, we see that both the highest overall
 201 success rate and the highest success rate in the worst-off neighborhood are obtained with equation (4). Finally, we find
 202 that equation (4) reduces the spread of income compared to (1), as can be seen from Figure 1c.




Figure 1: Results of the reproduction study. In subplots 1a and 1b, each marker represents an objective function. As in the original paper, success rates are small, because the number of drivers is much smaller than the number of riders. In subplot 1c, the income by objective function is shown.

203 4.2 Results beyond original paper

204 In addition to reproducing the results in the original paper, we performed supplementary experiments to test the
 205 generalizability of the methods used in the original paper. In subsection 4.2.1 we discuss the results of replacing the
 206 neural value estimator with a greedy value estimator. Further, in subsection 4.2.2 we examine the results of using our
 207 preprocessing method. Lastly, in subsection 4.2.3 we discuss the results obtained using the Green Taxi dataset.

208 4.2.1 Results using a greedy non-neural value estimator

209 These results are obtained by replacing the neural non-myopic value estimator with a myopic greedy value estimator.
 210 This serves both to test the added value of the neural estimator, and to test the sensitivity of the claims to a different
 211 estimation method. Interestingly, the results obtained with the greedy estimator (Figure 2) closely resemble the results
 212 obtained with the neural estimator (Figure 1), even in the case of 200 drivers where the neural estimator does train
 213 (as opposed to the case with 50 drivers). This suggests that, given the limited training examples for 200 drivers and 3
 214 days of taxi data, the neural model provides little added benefit over a myopic greedy estimator. By extension, this
 215 experiment reaches the same conclusions: claims 1 and 3 are replicated, but claim 2 is not.




Figure 2: Results of using a greedy value estimator. Apart from using a different value estimator, this experiment uses the exact same setup as the original paper (and that we use in our reproduction in section 4.1).

216 **4.2.2 Results using our preprocessing method**

217 To examine the sensitivity to changes in the preprocessing method, we compare the results of this experiment to the
 218 results of the reproduction. The obtained results, as displayed in Figure 3, are similar to the reproduced results. We see
 219 the same monotonic increase in Figures 3a and 3b as we saw in Figure 1 for the reproduction. Similarly, the income
 220 distribution (Figure 3c) shows that objective (4) obtains the highest income for the lowest-earning drivers and reduces
 221 the spread of income. Therefore, we can conclude that these results support claim 1 and claim 3, but do not support
 222 claim 2, and that the method proposed by the authors is robust to changes in the preprocessing method.




Figure 3: Results of using our preprocessing method. Apart from preprocessing the raw Yellow Taxi data differently, this experiment uses the exact same setup as the original paper (and that we use in our reproduction in section 4.1).

223 **4.2.3 Results on Green Taxi dataset**

224 In order to analyze the generalizability of the original paper, we use the results obtained on the Green Taxi dataset to see
 225 if the claims described in Section 2 hold for a different dataset. These results (Figure 4) are similar to the ones obtained
 226 on the Yellow Taxi dataset in both the reproduction experiment (Figure 1) and the experiment using our preprocessing
 227 (Figure 3). This experiment again supports claims 1 and 3, but does not support claim 2.

228 Important to note is that the success rates are much higher compared to those obtained on the Yellow Taxi dataset. This
 229 is because the Green Taxi demand in Brooklyn is significantly less than the Yellow Taxi demand in Manhattan. Hence,
 230 with an equal number of taxis a greater share of requests can be met. The fact that the findings of this experiment are
 231 consistent with the previous experiments, all of which have much lower success rates, provides confidence that the
 232 claims also generalize to realistic success rates (i.e. success rates approaching 100%, as is expected from real taxi
 233 companies).




Figure 4: Results for the Green Taxi Dataset. In subplots 4a and 4b, each marker represents an objective function. These results were obtained by averaging the success rates obtained for the two evaluation days. In subplot 4c, the income by objective function is shown. Here, the drivers from each evaluation day are treated separately.

234 5 Discussion

235 To conclude, our results support the first claim, since we find that the success rate in the worst-off neighborhood
236 increases monotonically with respect to the overall success rate for both 50 and 200 drivers. Unlike the original paper,
237 we do not find that the request-maximizing objective function (equation 1) maximizes income for the lowest-earning
238 drivers, so we do not support the second claim. Instead, we find that at 200 drivers the driver-side fairness objective
239 function (equation 4) obtains the highest income for the lowest-earning drivers.

240 We support the third claim since our results show that the driver-side fairness objective function yields the greatest
241 overall success rate, the highest success rate in the worst-off neighborhood, and reduces the spread of driver income.
242 Indeed, our results provide even stronger support for this claim than the original paper. Unlike the paper, we have found
243 the driver-side objective to provide the best success rates in both cases of 50 and 200 drivers. Furthermore, this objective
244 lifts the income of the least-earning drivers above what they make under the request-maximizing objective. This makes
245 the driver-fairness objective attractive in all respects, and an interesting subject for further research. One avenue to
246 explore is *why* the driver-fairness objective produces greater success rates than the request-maximizing objective, even
247 though the latter’s sole purpose is to maximize the success rate.

248 The obtained results are generally consistent across the four conducted experiments, showing that the claims which are
249 supported by our results (claims 1 and 3) are robust and relatively insensitive to a range of reasonable changes to the
250 experimental setup. This provides confidence that these claims and the corresponding objective functions generalize
251 well beyond the precise setups in which we and [4] tested them.

252 One limitation of our work, which is also shared by the original paper, is that the success rates across all experiments
253 are unrealistically low, because the number of drivers (50 or 200) is insufficient to meet demand (there are more
254 than 10,000 Yellow taxis in Manhattan). This creates an abundance of possible actions for each driver that is not
255 representative of the competition that exists for real-world taxi services. Running experiments in a setup where success
256 rates are more realistic would be a worthwhile additional generalization experiment. Such experiments may require
257 more computational resources than some researchers, ourselves included, have access to. However, our results for the
258 Green Taxi dataset already mitigate concerns that the claims would not generalize to greater success rates; the claims
259 were upheld in this setup with success rates of over 50%.

260 5.1 What was easy

261 The paper is written engagingly and the theoretical sections in particular give a clear description of the problem setup
262 and objectives. The paper is also accompanied by an open-source codebase with their implementation, which is
263 extremely helpful to obtaining accurate reproductions.

264 5.2 What was difficult

265 Even though the code was sufficient to reproduce the experiments done in the original paper, it was cluttered at times. It
266 contains functions that are never used, as well as print statements solely used to check if a certain point in the code is
267 reached without errors. Additionally, the code used to create the location embeddings creates embeddings of size 10,
268 when embeddings of size 100 are expected by the model. When reimplementing the code in PyTorch, most difficulty
269 was experienced when having to feed masked data into an LSTM backwards, as a result of how masking is implemented
270 in PyTorch. Further, various important details could only be found in the code, such as which days they used to train the
271 model on and the used epoch duration. One of the most important details that was missing is the notion that the model
272 does not train for 50 drivers and three days of Yellow Taxi data, as it will not train without enough examples. Further,
273 the data to perform the original experiments and the script to preprocess the raw data were not publicly available. After
274 contacting the authors, they provided us with the missing data. Despite what is stated in the paper, approximately 1% of
275 the routes in this data were not computed. These omitted routes, however, were not present in any of the rider requests
276 and therefore did not pose a problem to this research.

277 5.3 Communication with original authors

278 We contacted the authors because the data and the embeddings used to perform the original experiments are not publicly
279 available. The authors provided the missing data quickly and expressed willingness to help with further queries.

280 **References**

- 281 [1] New York City. “New York Yellow Taxi data, 2016”. In: ().
- 282 [2] Cody Cook et al. “The gender earnings gap in the gig economy: Evidence from over a million rideshare drivers”.
283 In: *The Review of Economic Studies* 88.5 (2021), pp. 2210–2238.
- 284 [3] Joanna Moody, Scott Middleton, and Jinhua Zhao. “Rider-to-rider discriminatory attitudes and ridesharing
285 behavior”. In: *Transportation Research Part F: Traffic Psychology and Behaviour* 62 (2019), pp. 258–273.
- 286 [4] Naveen Raman, Sanket Shah, and John Dickerson. “Data-Driven Methods for Balancing Fairness and Efficiency
287 in Ride-Pooling”. In: *arXiv:2110.03524 [cs]* (Oct. 7, 2021). arXiv: 2110.03524. URL: <http://arxiv.org/abs/2110.03524> (visited on 01/24/2022).
- 288
- 289 [5] Paolo Santi et al. “Quantifying the benefits of vehicle pooling with shareability networks”. In: *Proceedings of
290 the National Academy of Sciences* 111.37 (Sept. 16, 2014), pp. 13290–13294. ISSN: 0027-8424, 1091-6490. DOI:
291 10.1073/pnas.1403657111. arXiv: 1310.2963. URL: <http://arxiv.org/abs/1310.2963> (visited on
292 01/24/2022).
- 293 [6] Sanket Shah, Meghna Lowalekar, and Pradeep Varakantham. “Neural Approximate Dynamic Programming
294 for On-Demand Ride-Pooling”. In: *arXiv:1911.08842 [cs, eess]* (Nov. 20, 2019). arXiv: 1911.08842. URL:
295 <http://arxiv.org/abs/1911.08842> (visited on 01/24/2022).

Algorithm 1 NeurADP(N,T)

```

1: Initialize: replay memory  $M$ , Neural value function  $V$  (with random weights  $\theta$ )
2: for each episode  $1 \leq n \leq N$  do
3:   Initialize the state  $s_0^n$  by randomly positioning vehicles.
4:   Choose a sample path  $\xi^n$ 
5:   for each step  $0 \leq t \leq T$  do
6:     Compute the feasible action set  $F_t$  based on  $s_t^n$ .
7:     Solve the ILP in Table 1 to get best action  $a_t^n$ . Add the Gaussian noise for exploration.
8:     Store  $(r_t^n, F_t)$  as an experience in  $M$ .
9:     if t % updateFrequency == 0 then
10:      Sample a random mini-batch of experiences from  $M$ 
11:      for each experience  $e$  do
12:        Solve the ILP in Table 1 with the information from experience  $e$  to get the objective value  $y^e$ 
13:        for each vehicle  $i$  do
14:          Perform a gradient descent step on  $(y^{e,i} - V(r_t^{i,n}))^2$  with respect to the network parameters  $\theta$ 
15:        end for
16:      end for
17:    end if
18:    Update:  $s_t^{a,n} = T^a(s_t^n, a_t^n), s_{t+1}^n = T^\xi(s_t^{a,n}, \xi_{t+1}^n)$ 
19:  end for
20: end for

```

297 **B Results of the original paper**

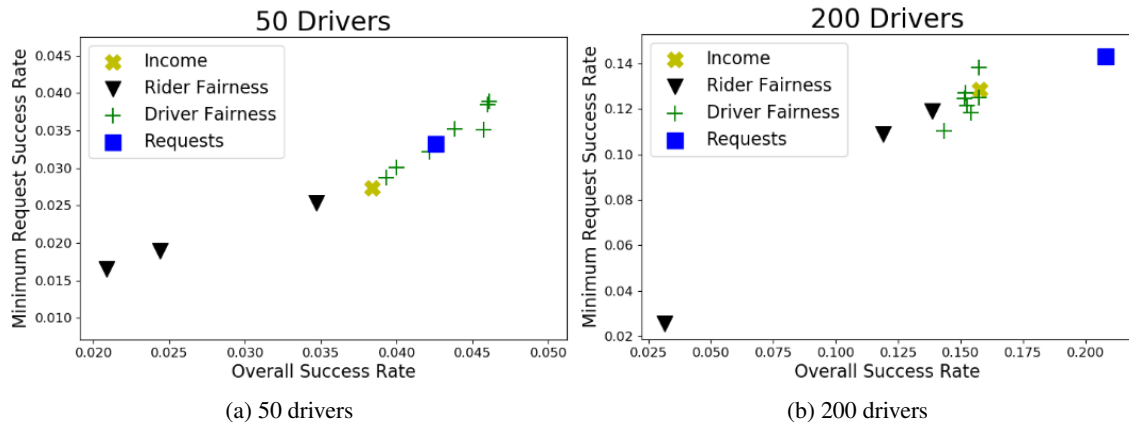


Figure 5: The results of the original paper [4]. The overall success rate is the rate of accepted requests in all neighborhoods, and the minimum request success rate is the minimum rate of accepted requests in a neighborhood. Each data point of the same symbol refers to a different (unspecified) value of λ .




Figure 6: Results of the original paper [4]. The income distribution for 200 drivers.