Optimization of the Revenue of the New York City Taxi Service using Markov Decision Processes

Jacky P.K. Li, Sandjai Bhulai

Vrije Universiteit Amsterdam,
Faculty of Science,
Amsterdam, The Netherlands
Email: jacky.li@vu.nl, s.bhulai@vu.nl

Theresia van Essen

Delft Institute of Applied Mathematics,
Delft, The Netherlands
Email: J.T.vanEssen@tudelft.nl

***Abstract*—Taxis are an essential component of the transportation system in most urban centers. The ability to optimize the efficiency of routing represents an opportunity to increase revenues for taxi drivers. The vacant taxis cruising on the roads are not only wasting fuel consumption, the time of a taxi driver, and create unnecessary carbon emissions but also generate additional traffic in the city. In this paper, we use Markov Decision Processes to optimize the revenues of taxi drivers by better routing. We present a case study with New York City Taxi data with several experimental evaluations of our model. We achieve approximately 10% improvement in efficiency using data from the month of January. The results also provide a better understanding of the several different time shifts. These data may have important implications in the field of self-driving vehicles.**

Keywords–New York taxi service; revenue optimization; optimal routing; Markov decision processes

I. Introduction

In New York City, there are over 485,000 passengers taking taxis per day, equating to over 175 million trips per year [1]. Creating an efficient way to transport passengers through the city is of utmost importance. Taxi drivers cannot control a passenger's destination but can make better decisions using optimal routing. This consequently leads to reductions of costs and of the carbon emissions.

Previous studies have focused on developing recommendation systems for taxi drivers [2]–[7]. Several studies use the GPS system to create recommendations for both the drivers and the passengers to increase the profit margin and cutting the time for seeking [4], [6]–[8]. Ge et al. [9] and Ziebart et al. [10] gather a variety of information to generate a behavior model to improve driving predictions. Ge et al. [2] and Tseng et al. [11] measure the energy consumption before finding the next passenger. Castro et al. [8], Altshuler et al. [12], Chawla et al. [13], Huang et al. [14] and Qian et al. [15] learn knowledge from taxi data for other types of recommendation scenarios such as fast routing, ride-sharing, or fair recommendations.

Most of the papers above focus on optimizing the measures for the immediate next trip. Rong et al. [3] investigate how to learn business strategies from the historical data to increase revenues of the taxi drivers using Markov decision processes (MDPs). Their research model uses historical data to estimate the probability of finding a passenger and its location for drop-off as the necessary parameters for the MDP model. For each one-hour time slot, the model learns a different set of parameters for the MDP from the data and finds the optimal move for the vacant taxi to maximize the total revenues in that time slot. At each state, the MDP model uses a combination of location, time, the current and the previous actions. The vacant taxi can travel to its neighboring locations and cruise through the grid to seek for the next passenger. Using dynamic programming to solve the MDP, the output of the model recommends the best actions for the taxi driver to take at each state.

Tseng et al. [11] examine the viability of electric taxis in New York City by using MDPs. Due to the usage limitation of electric taxis before each charge, they examine the profitability of replacing taxi with internal combustion engines by electric taxis. The research model uses OpenStreetMap (OSM) to assign each pick-up and drop-off into the nearest junctions. The advantage of using OSM is to be able to identify the number of available taxis at the junction without extra calculations. The research is concentrated on energy consumption; the actions become infeasible if the electric vehicle runs out of battery.

Analysis of real taxi data shows that there are significant differences in demand between certain periods of the day. The aforementioned research has not taken the effect of this demand variation into account. The contribution of our model is that we extend the research by Rong et al. [3] in this direction. We analyze the New York City Taxi data and study the differences in optimal policies and revenues for the demand between weekdays, weekends, day shifts, and night shifts. From these observations, we can infer relevant policies for taxi drivers based on the shift that they work in.

The paper is structured as follows. In Section II, we do data analysis on the New York Taxi dataset. This provides input for our MDP, which is explained in Section III. We assess the performance of the MDP in Section IV, where we conduct numerical experiments. Finally, the paper is concluded in Section V.

II. Dataset and Methodology

In our research, we use 14,776,615 taxi rides collected in New York City over a period of one month (January 2013) [1]. From each ride record, we use the following fields: taxi

 ISBN: 978-1-61208-603-3

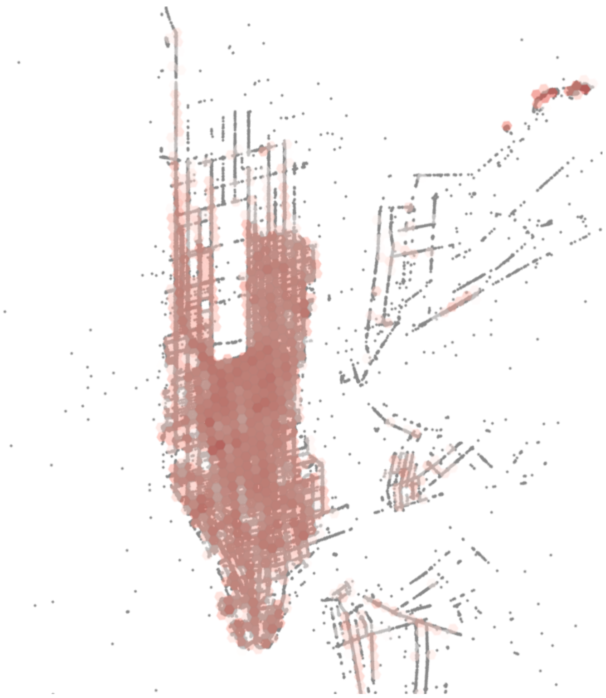

Figure 1. Rotated Manhattan with the total revenue for NYC Taxi by pick-up location in Jan 2013.

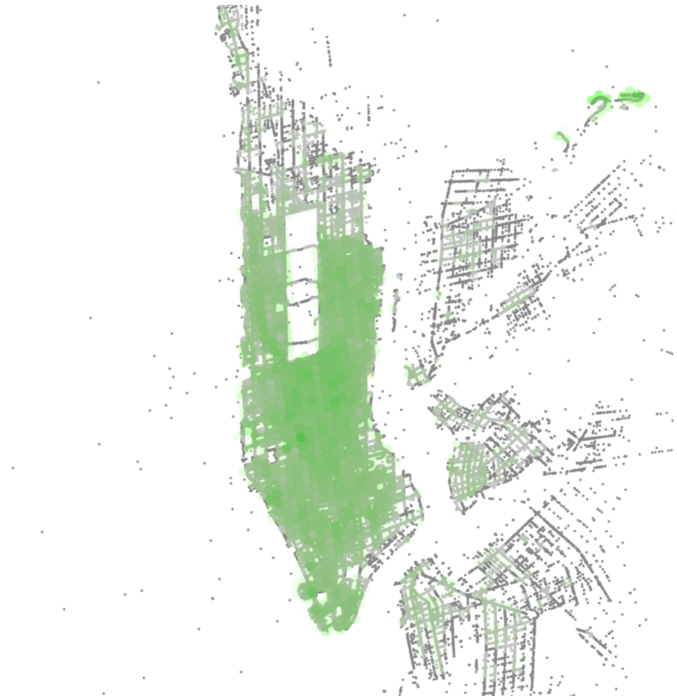

Figure 2. Rotated Manhattan with the total revenue for NYC Taxi by drop-off location in Jan 2013.

ID, pick-up time, pick-up longitude, pick-up latitude, drop-off time, drop-off longitude, drop-off latitude, the number of passengers per ride, average velocity, trip distance, traveling time, and fare amount. We omit the records containing missing or erroneous GPS coordinates. Records that represent rides that started or ended outside Manhattan, as well as trip durations longer than 1 hour and trip distances greater than 50 kilometers are omitted as well. Furthermore, we collect the drivers who drive for six to nine hours consistently to yield a clean dataset containing approximately 13.5 millions taxi rides. We observe that most of the pick-up locations are in the Manhattan area.

We concentrate on the island of Manhattan area in NY. This area imposes a rectangular grid of avenues and streets. However, the citys avenues are not parallel to the true north and south. For that reason, we tilted the map by 28.899 degrees according to Petzold et al. [16]. This creates blocks with the same grid system in most areas. We discretize the grid into a 50×50 grid, making each block in the grid approximately 300 meters $\times$ 300 meters. The choice for a block size of 300 meters is based on the assumption that a taxi can traverse this distance within 1 minute. Figure 1 shows the total revenue for the taxis by the pick-up location with the rotated map. Figure 2 indicates the total revenues of the drop-off location, and it shows that Lower Manhattan, along with the airport are the largest revenue generators and the drop-off location has spread to the mid-Manhattan area and also Brooklyn area.

The state of a taxi can be described by two parameters: the current location $L = \{(1,1), \ldots, (50,50)\}$ grid and the current time, $T = \{1, \ldots, 60\}$. We will denote the system state in our MDP model as $s = (x, y, t)$, which we will elaborate on in Section III.

A. *Performance indicators*

In this section, we present performance indicators of the taxi drivers. This will be used in the MDP to optimize the routing decision of each taxi driver. Hence, the performance indicators will be dependent on the routing policy that is being applied by the taxi drivers. To improve readability, we drop the dependency on the policy in the notation and use it only in cases where it benefits clarity.

We calculate the total business time of each taxi driver per shift. The total business time (denoted as T_{bus}) is equal to the sum of the total occupancy time (T_{occupy}) and the total seeking time (T_{seek}):

$$T_{\text{bus}} = T_{\text{occupy}} + T_{\text{seek}}. \tag{1}$$

The total occupancy time, T_{occupy} is the sum of all the trip durations with passengers of a taxi per day. And the total seeking time (T_{seek}) is the time between each trip. Figure 3 depicts the overall T_{seek} and the graphs in which we distinguish between the weekday, weekend, day shift, and the night shift. Based on the data, we consider 90% of the seeking time is less than 20 minutes for the day shift and less than 25 minutes for the night shift. Therefore, we discount any seeking time that is over 30 minutes as we assume those are the breaks for the drivers.

Logically, the T_{bus} is approximately the same for each taxi driver. To increase the revenue, the taxi drivers aim to have the maximal T_{occupy} and the minimal of T_{seek}. We define the revenue efficiency (E_{rev}) metric as the revenue earned divided by the total taxi drivers business time. This is expressed as follows:

$$E_{\text{rev}} = \frac{M}{T_{\text{bus}}} = \frac{M}{T_{\text{occupy}} + T_{\text{seek}}}, \tag{2}$$

 ISBN: 978-1-61208-603-3

Table I. REVENUE EFFICIENCY E_{REV}.

	Weekday dayshift	Weekday nightshift	Weekend dayshift	Weekend nightshift	Overall
Top 10%	0.59203	0.62408	0.60111	0.64646	0.60869
Mean	0.49985	0.52232	0.50252	0.54871	0.50565
Standard Deviation	0.07253	0.08011	0.07787	0.07799	0.08088
Bottom 10%	0.41028	0.42174	0.40426	0.44978	0.40572

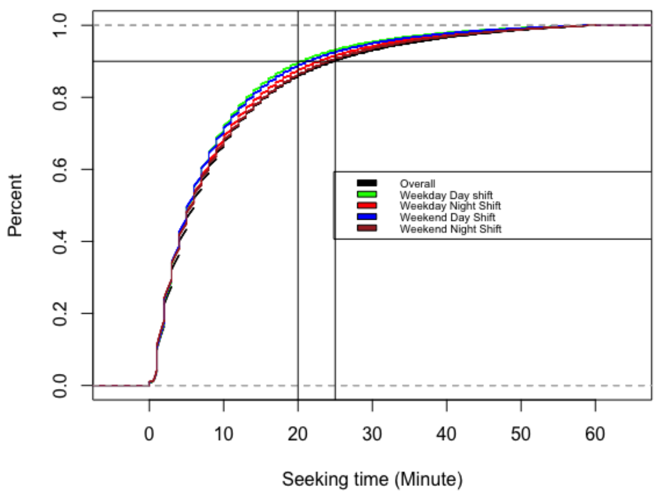

Figure 3. Seeking time for the models.

where M denotes the total money earned by the taxi driver during that period.

To illustrate the consistency of the taxi driver, we concentrate on the drivers who work between six hours to nine hours during the month of January. From that data, we generate the data of P_{find}, P_{dest}, T_{drive}, r (parameters of our MDP to be described in the next section) of each model and identify the top 10% and bottom 10% drivers in each model.

Table I indicates the revenue efficiency of the top 10% and bottom 10% distinguished by weekday, weekend, day shift, night shift, and the overall efficiency. Based on the table, there is approximately 20% difference between the performance of the top 10% and bottom 10% drivers. The previous studies that were mentioned above (see, e.g., [4], [8], [11], [13], [14]) attribute the difference between the performance by the top and bottom 10% drivers to the seeking time of the taxi drivers. This warrants research to determine if our model can provide a better solution for the taxi drivers for seeking passengers.

III. MARKOV DECISION PROCESS

In order to model the taxi service in New York City, we adopt the framework of MDPs. This framework allows us to deal with the uncertain demand over the different periods in the grid, and to model them explicitly. The MDP is a stochastic decision process with a set of states (S) and a set of possible actions (A) that transition the states from one to another. Each action will correspond to the process of the current state to the new state with a probability transition function and a reward function. The collection of optimal actions for each state is called the policy, which maximizes the total reward over several numbers of steps. The objective of our model is to minimize the seeking time for the taxi to maximize the expected revenues.

A. System States

The state for a taxi is described by its current locations and the current time. The details are explained as follows.

Location $(x, y) \in L = \{1, \ldots, 50 \times 1, \ldots, 50\}$: the area is divided into grid 50×50 grid cells;

Time $t \in T = \{1, \ldots, 60\}$: we use minutes as the interval of a time slot, and a total of 1 hour as time horizon.

Each pick-up and drop-off location is assigned to a grid cell. We remove the records that contain $1)$ incomplete data information, $2)$ trip distance over 100 kilometers, $3)$ trip durations over 60 minutes, $4)$ pick-up and drop-off locations with the same coordinates, $5)$ pick-up and drop-off locations outside the grid, and $6)$ shifts that are shorter than six hours and longer than nine hours.

We denote the system state of our MDP model as $s = (x, y, t)$, and the collection of all admissible states is denoted by S.

B. Actions

The admissible actions from a given state s have nine possibilities to choose from. We use numbers $1, \ldots, 9$ to index the directions. We express them formally as:

$$A = \begin{array}{|c|c|c|}\hline 7 & 8 & 9 \\ \hline 4 & 5 & 6 \\ \hline 1 & 2 & 3 \\ \hline\end{array},$$

where, e.g., action 9 moves the taxi to the neighboring north-east location.

C. Parameters of the MDP model

In this subsection, we state the parameters used in the rest of MDP model.

The probability parameters are defined as:

- $P_{\text{find}}\,(x, y)$ describes the probability of successfully picking up a passenger in grid cell (x, y). We can calculate the probability of picking up a passenger in the cell by dividing the number of successful pick-ups in the cell $(n_{\text{find}}(x, y))$ by the total number of times this cell is visited by a vacant taxi. The vacant taxi includes the taxis that drop off passengers in

Copyright (c) IARIA, 2017. ISBN: 978-1-61208-603-3

grid cell (x, y) ($n_{\text{drop-off}}(x, y)$) and also the taxis that are seeking for passengers ($n_{\text{OSRM}}(x, y)$). To locate the vacant taxi every minute during the seeking trip, we use the API provided by Open Source Routing Machine [17], to estimate the coordinates. We use one hour time slots between 12:00 to 13:00 for the day shift model and 0:00 to 1:00 for the night shift model. In our overall model, we took the average of the day time and night time models to estimate the number of vacant taxis at each grid during the month of January in 2013. Thus,

$$P_{\text{find}} = \frac{n_{\text{find}}(x, y)}{n_{\text{find}}(x, y) + n_{\text{drop-off}}(x, y) + n_{\text{OSRM}}(x, y)}.$$

- $P_{\text{dest}}(x, y, x', y')$ describes the probability of a passenger travelling from grid cell (x, y) to the grid cell (x', y'). To estimate the destination probability for a time slot, we calculate the number of trips between each pair of source and destination locations in that time slot and get a 50×50 matrix. The value is divided by the sum of the entire number of trips of the grid cells. Therefore, P_{dest} has the empirical probability distribution of a passenger choosing destination location (x', y') when he is picked up at location (x, y).

The time parameters are defined as:

- $T_{\text{seek}}(a)$: The required time to travel from one location to a neighboring location based on action $a \in A$. We assume that the average speed of seeking trips is approximately 300 meters per minute. Thus, a taxi can traverse on cell when $a = 2, 4, 5, 6, 8$, and hence $T_{\text{seek}}(a) = 1$ in this case. In case $a = 1, 3, 7, 9$, then we set $T_{\text{seek}}(a)$ equal to 2, due to the diagonal movement.
- $T_{\text{drive}}(x, y, x', y')$: The driving time from (x, y) to (x', y'). We can calculate the total driving time from grid cell (x, y) to grid cell (x', y') and then divide by the number of trips from grid cell (x, y) to grid cell (x', y'). We calculate T_{drive} individually for all models. From the calculation, there is approximately +15.67% driving time difference between the day shift model and the night shift model, and there is a +4.14% difference between the weekend and the weekday.
- We assume there is no waiting time for passengers to get in and out of the vehicle.

The reward is defined as:

- $r(x, y, x', y')$: The expected reward from grid cell (x, y) to grid cell (x', y'). Similar to T_{drive}, we calculate the average fare of the number of trips between each pair of source and destinations as the expected fare. Note that due to this definition, we reward does not depend on the action of the taxi driver. We calculate r separately for all models. Similarly to T_{drive}, there is approximately +6.21% reward difference between the day shift model and the night shift model, and there is a +1.21% difference between the weekend and the weekday.

D. State transition function

The state transition function is a function that describes the probability that one moves from state (x, y, t) after taking decision a moves to state (x', y', t'). Assuming the current state is $S = (x, y, t)$ and action a is taken, there are two possible outcomes of the transition:

1) The taxi successfully finds a passenger in grid (x, y) within $T_{\text{seek}}(a)$ minutes. The taxi with the passenger goes to destination (x', y') with probability $P_{\text{dest}}(x, y, x', y')$. The taxi arrives at location (x', y') with $T_{\text{drive}}(x, y, x', y')$ as the total time used to travel from (x, y) to (x', y'). The taxi driver receives $r(x, y, x', y')$ as the expected reward. Then the taxi will start seeking for a passenger from grid cell (x', y'). In this case, the new state becomes $s' = (x', y', t + T_{\text{seek}}(a) + T_{\text{drive}}(x, y, x'y'))$.
2) The taxi does not find a passenger after $T_{\text{seek}}(a)$ minutes being in grid (x, y) with the probability $(1 - P_{\text{find}}(x, y))$. The taxi driver does not receive a reward and saves the driving time T_{drive}. The taxi driver starts to make the next action at grid cell (x', y'). Hence, the state of the taxi driver becomes $s' = (x', y', t + T_{\text{seek}}(a))$.

E. The objective function

The objective function of the MDP model is to maximize the total expected rewards starting from an initial state. The terminal states are the states with $t = 60$. No more actions can be taken once the system reaches the terminal states. The maximal expected reward for an action a in state $s = (x, y, t)$ is expressed as $V(s, a)$ shown in (3).

$$\begin{aligned} V(s,a) &= (1 - P_{\text{find}}(x,y)) \times \\ &\max_{a' \in A} V(x, y, t + T_{\text{seek}}(a), a') + \\ &\sum_{(x',y') \in L} P_{\text{find}}(x,y) \times P_{\text{dest}}(x,y,x',y') \times \\ &\big[r(x,y,x',y') + \\ &\max_{a' \in A} V(x', y', t + T_{\text{seek}}(a) + T_{\text{drive}}(x,y,x',y'), a') \big]. \end{aligned} \tag{3}$$

The optimal policy π^* is defined as:

$$\pi^*(s) = \arg\max \{V(s, a)\}, \tag{4}$$

and the optimal value function is given by

$$V^*(s) = V(s, \pi^*(s)). \tag{5}$$

F. Markov Decision Process Solution

In order to solve the Markov decision problem to derive the optimal policy, we employ dynamic programming to maximize the expected rewards. The algorithm starts from time $t = 60$ and then traces backward to time $t = 1$. The algorithm is listed in Algoritm 1.

 ISBN: 978-1-61208-603-3

Table II. REVENUE EFFICIENCY E_{REV}.

	Weekday dayshift	Weekday nightshift	Weekend dayshift	Weekend nightshift	Overall
Top 10%	0.59203	0.62408	0.60111	0.64646	0.60869
$P_{\text{find}}(x, y)$	0.52267	0.50915	0.51463	0.45475	0.50030
Bottom 10%	0.41028	0.42174	0.40426	0.44978	0.40572

Algorithm 1 Solving MDP using Dynamic Programming

Input: $L, A, T, P_{\text{find}}, P_{\text{dest}}, r, T_{\text{drive}}, T_{\text{seek}}$
Output: The best policy π^*

V is a $\mid L \mid \times \mid T \mid$ matrix; $V \leftarrow 0$
for $t = \mid T \mid$ to 1 do
 for all $(x, y) \in L$ do $\triangleright s = (x, y, t)$
 $a_{max} \leftarrow a$ that maximizes $V(s, a)$
$\pi^*(s) \leftarrow a_{max}$
$V^*(s) \leftarrow V(s, a_{max})$
return π^*

IV. CASE STUDY

In this section, we present our case study on the New York Taxi dataset. We evaluate the MDP for the expected reward based on the dataset from January 2013. We assume that the NYC taxis have two shifts per day and each shift is a 12-hour period. We analyze the taxi's expected reward in 1) the day-time shift within six to nine hours of its operating time, 5 am to 5 pm and 2) the night-time shift, 5 pm to 5 am and 3) the weekdays from Monday to Friday, and 4) the weekend from Friday to Sunday. After filtering the data, we have approximately 170,000, 205,000, 145,000, and 193,000 shifts, respectively, for the Weekday day-time shift, Weekday night-time shift, Weekend day-time shift, and Weekend night-time shift. Although the weekend has a fewer number of days in January, the total number of shifts of the weekend night time is almost the same as for the weekday night time.

The results of the case study (see also Table II) shows that in our model

- $P_{\text{find}}(x, y)$ is 0.52267 which is 27.58% better than the bottom 10%, and it is 11.65% less effective than the top 10% for the Weekday day-time model.
- For the weekday night-time model, $P_{\text{find}}(x, y)$ is 0.50915 which is 27.52% better than the bottom 10%. It is 16.74% less effective than the top 10%.
- For the weekend day-time model, $P_{\text{find}}(x, y)$ is 0.51463 which is 20.16% better than the bottom 10%. It is 18.57% less effective than top 10%.
- For the weekend nighttime model, $P_{\text{find}}(x, y)$ is 0.45475 which is almost the same as the bottom 10% and it is 29.14% less effective than the top 10%.
- The overall model, $P_{\text{find}}(x, y)$ is 0.50030 which is 23.41% better than the bottom 10% and it is 17.79% less effective than top 10%.

The results of the case study show that our model is capable of reducing the time to find a passenger for a taxi driver significantly. Consequently, the end result is that the earnings of the taxi drivers increases. This benefit is expressed as approximately a 10% improvement in efficiency.

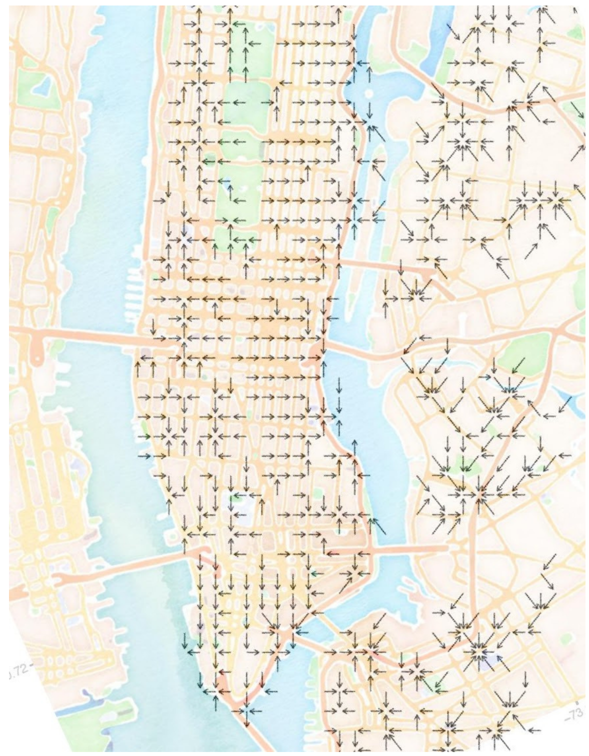

Figure 4. Recommended movements by the MDP model.

V. CONCLUSION AND FUTURE DISCUSSION

In this paper, we use MDP to model the taxi service strategy and determine the optimal policy for taxi drivers with a daytime and nighttime model during the weekdays and the weekend. This paper proposed to model the passenger-seeking process to receive the best move for a taxi that is seeking for the next passenger. Figure 4 shows the recommended movement by the MDP Model.

From the results of the case study, we observe that the weekend night time draws an interesting discussion. It has a similar number of shifts as compared to the weekday night time model, but the revenue efficiency did not improve compared to the bottom 10% drivers. A possible explanation might be that the experienced drivers would use their experience to look for the best location to seek customers. Consequently, the data may not have provided enough evidence to improve the bottom 10% drivers.

In our data analysis, we found cases where there are pick-up and drop-off locations in the Hudson River. We can assume that this is an error in the GPS system. Similar to this issue,

 ISBN: 978-1-61208-603-3

P_{dest} was estimated from a small number of trips from one location to another. This could sometimes result in a high probability, for instance, 1 of 3, would have created a 33% of probability going from one location to another. Further research is needed to develop methods to get a more accurate estimate.

REFERENCES

[1] N. Taxi, L. Commission et al., "2014 taxicab fact book," 2014.

[2] Y. Ge et al., "An energy-efficient mobile recommender system," in Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining - KDD '10. New York, New York, USA: ACM Press, 2010, p. 899. [Online]. Available: http://dl.acm.org/citation.cfm?doid=1835804.1835918

[3] H. Rong, X. Zhou, C. Yang, Z. Shafiq, and A. Liu, "The rich and the poor: A markov decision process approach to optimizing taxi driver revenue efficiency," in Proceedings of the 25th ACM International on Conference on Information and Knowledge Management. ACM, 2016, pp. 2329–2334.

[4] J. Yuan, Y. Zheng, L. Zhang, X. Xie, and G. Sun, "Where to find my next passenger," in Proceedings of the 13th international conference on Ubiquitous computing. ACM, 2011, pp. 109–118.

[5] Y. Zheng, J. Yuan, W. Xie, X. Xie, and G. Sun, "Drive Smartly as a Taxi Driver," in 2010 7th International Conference on Ubiquitous Intelligence & Computing and 7th International Conference on Autonomic & Trusted Computing. IEEE, oct 2010, pp. 484–486. [Online]. Available: http://ieeexplore.ieee.org/document/5667121/

[6] M. Qu, H. Zhu, J. Liu, G. Liu, and H. Xiong, "A cost-effective recommender system for taxi drivers," in Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining - KDD '14. New York, New York, USA: ACM Press, 2014, pp. 45–54. [Online]. Available: http://dl.acm.org/citation.cfm?doid=2623330.2623668

[7] D. Zhang et al., "Understanding Taxi Service Strategies From Taxi GPS Traces," IEEE Transactions on Intelligent Transportation Systems, vol. 16, no. 1, feb 2015, pp. 123–135. [Online]. Available: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6841047

[8] P. Castro, D. Zhang, and S. Li, "Urban traffic modelling and prediction using large scale taxi gps traces," Pervasive Computing, 2012, pp. 57–72.

[9] Y. Ge, C. Liu, H. Xiong, and J. Chen, "A taxi business intelligence system," in Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining - KDD '11. New York, New York, USA: ACM Press, 2011, p. 735. [Online]. Available: http://dl.acm.org/citation.cfm?doid=2020408.2020523

[10] B. D. Ziebart, A. L. Maas, A. K. Dey, and J. A. Bagnell, "Navigate like a cabbie: Probabilistic reasoning from observed context-aware behavior," in Proceedings of the 10th international conference on Ubiquitous computing. ACM, 2008, pp. 322–331.

[11] C.-M. Tseng and C.-K. Chau, "Viability analysis of electric taxis using new york city dataset," in Proceedings of the Eighth International Conference on Future Energy Systems. ACM, 2017, pp. 328–333.

[12] T. Altshuler, R. Katoshevski, and Y. Shiftan, "Ride sharing and dynamic networks analysis," arXiv preprint arXiv:1706.00581, 2017.

[13] S. Chawla, Y. Zheng, and J. Hu, "Inferring the Root Cause in Road Traffic Anomalies," in 2012 IEEE 12th International Conference on Data Mining. IEEE, dec 2012, pp. 141–150. [Online]. Available: http://ieeexplore.ieee.org/document/6413908/

[14] Y. Huang, F. Bastani, R. Jin, and X. S. Wang, "Large scale real-time ridesharing with service guarantee on road networks," Proceedings of the VLDB Endowment, vol. 7, no. 14, 2014, pp. 2017–2028.

[15] S. Qian, J. Cao, F. L. Mouël, I. Sahel, and M. Li, "Scram: a sharing considered route assignment mechanism for fair taxi route recommendations," in Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. ACM, 2015, pp. 955–964.

[16] C. Petzold. How far from true north are the avenues of manhattan? [Online]. Available: http://www.charlespetzold.com/etc/AvenuesOfManhattan/ (2015)

[17] D. Luxen and C. Vetter, "Real-time routing with openstreetmap data," in Proceedings of the 19th ACM SIGSPATIAL international conference on advances in geographic information systems. ACM, 2011, pp. 513–516.

 ISBN: 978-1-61208-603-3