Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Computing and Information Systems | School of Computing and Information Systems

8-2021

Automated taxi queue management at high-demand venues

Mengyu JI
Singapore Management University, mengyuji.2017@phdis.smu.edu.sg

Shih-Fen CHENG
Singapore Management University, sfcheng@smu.edu.sg

Follow this and additional works at: https://ink.library.smu.edu.sg/sis_research

Part of the Artificial Intelligence and Robotics Commons, and the Computer Engineering Commons

Citation

1

This Conference Proceeding Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylds@smu.edu.sg.

2021 IEEE 17th International Conference on Automation Science and Engineering (CASE)
August 23-27, 2021. Lyon, France

Automated Taxi Queue Management at High-Demand Venues

Mengyu Ji[1] and Shih-Fen Cheng[2]

Abstract— **In this paper, we seek to identify an effective management policy that could reduce supply-demand gaps at taxi queues serving high-density locations where demand surges frequently happen. Unlike current industry practice, which relies on broadcasting to attract taxis to come and serve the queue, we propose more proactive and adaptive approaches to handle demand surges. Our design objective is to reduce the cumulative supply-demand gaps as much as we could by sending notifications to individual taxis. To address this problem, we first propose a highly effective passenger demand prediction system that is based on the real-time flight arrival information. By monitoring cumulative passenger arrivals, and control for factors such as the flight's departure cities, we demonstrate that a simple linear regression model can accurately predict the number of passengers joining taxi queues. We then propose an optimal control strategy based on a Markov Decision Process to model the decisions of notifying individual taxis that are at different distances away from the airport. By using a real-world dataset, we demonstrate that an accurate passenger demand prediction system is crucial to the effectiveness of taxi queue management. In our numerical studies based on the real-world data, we observe that our proposed approach of optimal control with demand predictions outperforms the same control strategy, yet with Poisson demand assumption, by 43%. Against the status quo, which has no external control, we could reduce the gap by 23%. These results demonstrate that our proposed methodology has strong real-world potential.**

I. Introduction

In many major cities around the globe, taxi-like services (both traditional taxis and ride-hailing services such as Uber or Lyft) have become more and more important in providing point-to-point transportation services. This is especially true for locations that generate large amount of demands in a spiky manner and are relatively isolated. Best examples of such locations are airports, rail stations, stadiums, and conference/exhibition centers. It is common to see hundreds of people requesting for rides all at the same time during the demand peaks, and this usually results in significant waiting time for passengers.

A common approach adopted by most venue operators to address such issue is to broadcast potential needs for taxis (or ride-hailing cars) ahead of or during the demand peaks. However, most implementations of such approach is ad hoc in nature (in terms of the number of cars requested and the timing of the request), and is rarely coordinated across the planning horizon (when demand peak lasts for prolonged periods of time). As a result, a balance in passenger demand and vehicle supply for rides is rarely achieved. This recurring problem calls for new solution approaches, and our work in this paper is an attempt to address this issue.

In this paper, we propose an automated taxi queue (or ride-hailing cars) management framework that combines demand-side predictive analytics with supply-side policy optimization to come up with recommendations on the number of drivers to contact during different time periods, in anticipation of current and predicted future demand levels. Although the problem of taxi queue (or ride-hailing cars) management at high-demand locations is general in nature, and can be applied to many different types of venues, each venue type brings different set of challenges, which can be driven by passenger demand patterns, venue design, and geographical locations.

To concretely demonstrate the value of combining both predictive analytics and control policy optimization, we choose to focus on the setting of an airport, more specifically, the Changi Airport in Singapore. Before the COVID-19 pandemic, Changi Airport was consistently ranked as one of the best and busiest airports in the world. For an airport that handles more than 68 million passenger traffic (in 2019), it is extremely challenging to maintain a high level passenger experience. And the optimization of the taxi queue operations will be important in achieving this objective.

In solving the real-world use case at the airport, we aim to make the following contributions:

- We first propose a highly effective passenger demand prediction system that is based on the real-time flight arrival information. By monitoring *cumulative* passenger arrivals, and control for factors such as flight's departure cities, we demonstrate that a simple linear regression model can accurately predict the number of passengers joining taxi queues.
- We then propose an optimal control strategy based on a Markov Decision Process to model the decisions of notifying individual taxis that are at different distances away from the airport. By using a real-world dataset, we demonstrate that an accurate passenger demand prediction system is crucial to the effectiveness of taxi queue management. In our numerical studies based on the real-world data, we observe that our proposed approach of optimal control with demand predictions outperforms the same control strategy, yet with Poisson demand assumption, by 43%. Against the status quo, which has no external control, we could reduce the gap by 23%. These results demonstrate that our proposed methodology has strong real-world potential.

*This work was supported in part by the Changi Airport Group

[1]Mengyu Ji is currently a PhD student with School of Computing and Information Systems, Singapore Management University mengyuji.2017@phdis.smu.edu.sg

[2]Shih-Fen Cheng is currently an Associate Professor with School of Computing and Information Systems, Singapore Management University sfcheng@smu.edu.sg

II. Related Work

There are two major streams of research that are related to our research. The first stream of research focuses on demand and waiting time predictions at taxi queues. The second stream of research focuses on queue management.

A. Passenger Demand Predictions at Taxi Queues

As more and more microscopic data on taxi movement and trips become available, there have been a surge in research on taxi demand prediction, particularly ones that utilize advanced machine learning techniques. For example, Yao et al. [1] have introduced the use of deep multi-view neural network in taxi demand prediction, while Geng et al. [2] build upon similar deep neural network framework to incorporate contextual information from the traffic network. Although these techniques work well in practice, they aim to make prediction at the city level, and not designed to handle a few queues with high demands.

With a focus on queues, Lu et al. [3] have studied how to utilize taxi mobility traces to detect the formation of queues and the associated context (e.g., passengers or taxis). Rahaman et al. [4] have studied the prediction of taxi queue context in airports, using taxi driver's knowledge; they have further developed this line of research to also look at the imbalances of taxi and passenger demands at queues. Using modern machine learning techniques, Rahaman et al. [5] also look into using $k-NN$-based regression methods to predict taxi queue waiting time, and it allows a large number of external factors to be considered.

Compared to these past research, our contributions in this area are unique in the following ways:

- We focus on developing prediction models for the airport.
- Recognizing the uniqueness of the airport context, we incorporate the flight arrival information to significantly improve the prediction quality of passenger demands.

B. Taxi Queue Management

An early study by Curry et al. [6] study the interaction between taxicabs and buses at an airport, but their focus is on passenger waiting time in order to reduce airport ground transport congestion. In the past decade, as GPS-based and smartphone-based technology becoming more mature and accessible, taxi fleet operators are increasingly utilizing these technologies to manage their fleets, resulting in the availability of many large-scale microscopic datasets. These datasets have enabled researchers to conduct realistic studies of taxi queue management. For example, Cheng and Qu [7] propose a service choice model to help individual drivers in deciding whether to serve a specific taxi stand or not. Kamga et al. [8] study taxicab management at busy airports such as the JFK airport in the New York City, and reviewed how computerized taxicab dispatch system can be utilized to reduce passenger waiting time. More recently, Anwar et al. [9] propose a smartphone-based App to instruct taxi drivers to visit the most promising taxi queues at the Changi Airport in Singapore, taking into account average arrival time. In a broader city context, we have recently seen the implementation and field trial of a driver guidance system for taxi drivers [10], which is shown to help taxi drivers following guidance reduce their vacant roaming time by an average of 27% [11].

Compared to these past research, our contributions in this area are unique in the following ways:

- We choose to minimize the difference between demand and supply of taxis at a specific location, thus taking care of the welfare of both the drivers and the passengers.
- In the context of airport taxi queue management, we illustrate the importance of utilizing high-quality real-time demand prediction for taxi queue management. The resulting framework we propose integrates both taxi queue-specific prediction model and the guidance model, and we illustrate the superiority of such design.

III. Passenger Inflow Prediction

In most airports, the information on arriving passenger flights is known in advance through the flight information service, which can include (not exhaustively) the flight number, the origin, the final destination (if the current airport is a transit location), the number passengers on board, the breakdown of passenger types (e.g., arriving vs. transit, adult vs. children), the number of luggages, and the expected/actual arrival times at the gate. This information allows airport operators to acquire the exact number of passengers that will be flowing into the airport from different gates in real-time, and plan many important airport operations accordingly.

However, even with this passenger arrival information, predicting the number of passengers who will be joining a taxi queue at different time periods is still highly challenging, for the following reasons: 1) there are a wide variety of ground transport options available to passengers, thus it's uncertain how many of arriving passengers will eventually choose to join the taxi queue, 2) each flight arrives at different gate, resulting in different movement time, and 3) arriving passengers would need to clear immigration, pick up luggages (if any), and may be delayed along the way for many personal reasons. As a result, even for the same flight that arrives daily at roughly the same time, the number of passengers who would eventually show up at various time periods after this flight's arrival still varies greatly from day to day.

By using the actual flight arrival data and the taxi queue inflow data (collected via highly accurate LiDAR-based crowd sensors) at the Singapore Changi Airport. We confirm that it is indeed impractical to try to establish the relationship between individual flight's arrivals and the number of passengers that will end up joining the taxi queue.

However, the correlation becomes much more reliable after we focus on the **cumulative** flight arrivals and taxi queue inflows. To evaluate this alternative way of predicting taxi queue inflows, we designate accumulation time windows for both flight arrivals (observed) and taxi queue inflows (to be predicted) following the timeline illustrated in Figure 1.

In the initial analysis, we monitor the cumulative passenger counts from all flights that land within the 40-minute time window. This time window should end 10 minutes before the current time. The cumulative taxi queue inflows to be predicted should cover the next 15 minutes.

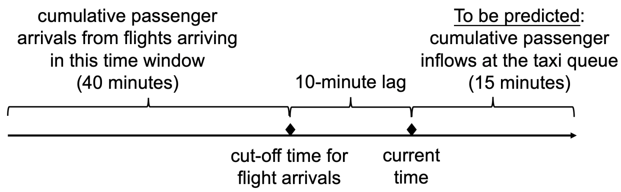

Fig. 1. Timeline for the cumulative arrival analysis.

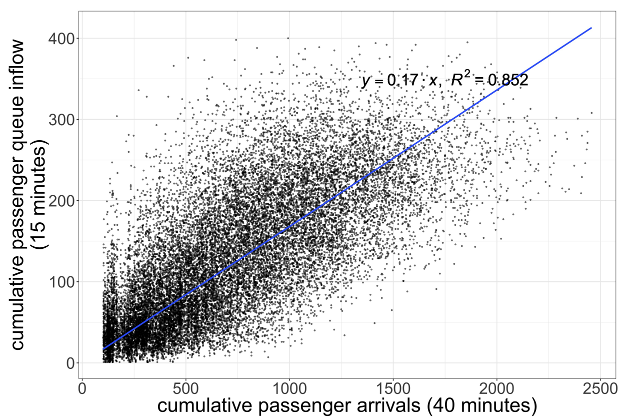

Fig. 2. Correlations between cumulative flight arrivals and taxi queue inflows.

The correlation between cumulative taxi queue inflows versus cumulative flight passenger arrivals is illustrated in Figure 2. With the $R^2 = 0.852$, we can see that even without any additional features, the linear correlation is strong. The linear regression formula suggests that around 17% of all passengers arrived during the 40-minute time window, after a 10-minute lag, would show up at the taxi queue during the 15-minute time window.

To further enhance the predictive power of this cumulative model, we incorporate a number of external parameters such as origin cities of arriving flights and the distance from the gate to the airport exit. The details on how we incorporate these additional features and the comparison of variants of our cumulative prediction model against other well-known prediction baselines are summarized next.

A. Cumulative Prediction Model

To account for the impact of the origin city and the gate of arrival for a flight, we structure the linear regression using the following formula:

$$f^{15} \sim \sum_{c \in \mathbf{C}, g \in \mathbf{G}} \beta^{40}_{c,g} \, x^{40}_{c,g},$$

where f^{15} denotes the taxi queue inflows to be predicted in the next 15 minutes, $x^{40}_{c,g}$ denotes the number of arriving passengers from city c and gate category g[1] within the 40-minute time window, and $\beta^{40}_{c,g}$ be the regression coefficient associated with $x^{40}_{c,g}$. As not all cities have significant enough passenger arrivals for the purpose of regression analysis, we aggregate them according to the countries or regions (e.g., Western/Eastern Europe, Greater China area, South-East Asia). For the gate category, we assign gates to 3 categories depending on the distance from the arrival gate to the airport exit.

The training dataset comes from October and November, 2019, while the testing dataset comes from December 2019. To see the benefits of having the city and gate features, we first execute the regression with city/gate features (see Table I) for all 3 terminals at the airport that we study; this set of experiments are then followed by the regression analysis without city/gate features (see Table II). From Tables I and II, we can see that the complete linear regression model with city/gate features performs better than the plain version of the model without city/gate features, by around 2% to 5% across different error measures. All coefficients in our complete linear regression model are statistically significant to 1%.

TABLE I
THE PERFORMANCE OF THE COMPLETE LINEAR REGRESSION MODEL WITH CITY/GATE FEATURES.

Terminal	RMSE	MAPE	SMAPE
Terminal 1	15.55	32.02%	26.43%
Terminal 2	12.70	48.35%	34.29%
Terminal 3	12.91	49.56%	35.89%

TABLE II
THE PERFORMANCE OF THE LINEAR REGRESSION MODEL WITHOUT CITY/GATE FEATURES.

Terminal	RMSE	MAPE	SMAPE
Terminal 1	16.22	32.89%	27.86%
Terminal 2	13.20	51.07%	35.66%
Terminal 3	13.25	51.06%	36.68%

The predictions against ground truth data for the testing dataset are plotted below in Figures 3 and 4.

B. Cumulative Prediction Model Against Baselines

To probe whether it is worthwhile to go beyond the linear regression model, we also compare our linear regression model against two other popular prediction baselines: ARIMA and XGBoost. The ARIMA model is identified as ARIMA(1,1,2) using the standard Auto-ARIMA forecast package. The XGBoost model is also based on standard implementation. We use the same dataset as before, and the comparison results are summarized in Table III.

From Table III, we can see that measured by SMAPE, our complete linear regression model outperforms ARIMA

[1]Depending on the walking distance, we classify all arrival gates into different categories based on the distance from the gate to the airport exit.

and XGBoost models by 9.7% and 2.3% respectively. This implies our linear regression model is sufficient for our prediction purpose.

TABLE III
LINEAR REGRESSION MODEL AGAINST BASELINES.

Model	RMSE	MAPE	SMAPE
ARIMA	17.00	41.44%	29.28%
XGBoost	15.77	33.77%	27.06%
Linear Regression	15.55	32.02%	26.43%

C. Multi-Period Prediction

Our linear regression model is designed to provide a single-period (15-minute) prediction for the inflow that will happen in the next 15 minutes. To extend our prediction to the next 30 minutes, we will hypothetically roll forward the current time by 15 minutes into the future, and update the value of independent variables to reflect the updated 40-minute time window for the flight arrivals.

IV. A MARKOV DECISION PROCESS FOR TAXI SUPPLY OPTIMIZATION

In this section, we propose a Markov decision process (MDP) to model how we can optimize the taxi supply at the taxi queues by sending personalized messages. This is in contrast to broadcasting to all taxis, which is simple yet reactive and not able to respond to current and future demand beforehand. In the below subsections, we define all important features of the MDP in greater details.

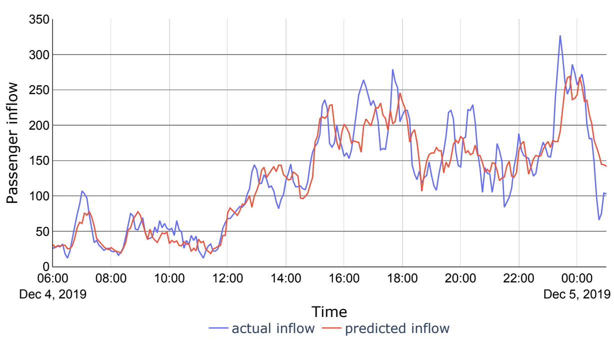

Fig. 3. Predicted values vs. ground truth: Dec 4, 2019

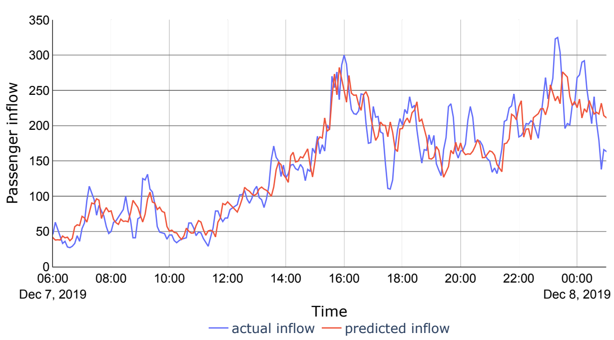

Fig. 4. Predicted values vs. ground truth: Dec 7, 2019

Our MDP formulation has discrete time periods, with the sequence of event occurrence specified in Figure 5. We assume that all actions from the controller (sending notifications) occur at the beginning of the period (the epoch). Drivers who receive notifications will decide on the spot whether to accept the notifications (and come to serve the taxi queue), or ignore them (and follow transition dynamics). The actual movement of taxis and the arrivals of passengers happen during the time period between two epochs. At the end of current period, and right before the start of the next epoch, all state variables (to be defined later) will be updated.

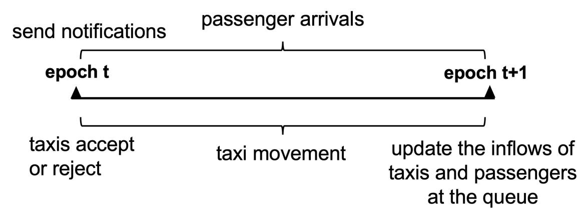

Fig. 5. Decision epochs and the event occurrence sequence in the MDP.

To make the MDP formulation more manageable, we define a small number of zones around the airport to track the taxi supply (instead of monitoring all taxis' current locations, which would be intractable in practice). In Figure 6, we provide an illustration of the zone definition and how transitions can occur between zones. The sizes of the zones are chosen such that a taxi can travel from any zone to adjacent zone in one time period.

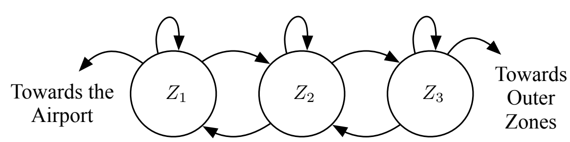

Fig. 6. Zones around airport.

In the absence of control, we assume that taxis would follow a Markov chain model to travel between zones. The transition probabilities between zones can be derived from the empirical dataset. For our case study, we have obtained a proprietary taxi movement dataset that contains taxi's periodic movement information, from which we can infer taxi's zone- and time-dependent transition probabilities to nearby zones. This will serve as the background movement model without control intervention.

A. The MDP Formulation

The optimal taxi queue control problem is formulated as a finite-horizon MDP, using the notations defined in Table IV. We denote the MDP as a tuple of $< S, A, T, R >$, where S is the state space, A is the action space, T is the transition probability, and R is the reward function. Both T and R are defined as functions, where $T : S \times A \times S \to [0, 1]$ and $R : S \times A \times S \to \mathbb{R}$.

At time t, the MDP state is defined as:

$$< supply_t^i,\ otw_t^i,\ W_{p_t}^j,\ W_{v_t}^j >, \forall i,\ j.$$

The action a_t^i is defined as the number of taxis we are going to engage for each zone i and each time period t, while the index j is use to denote the number of periods taxis or passengers have already spent waiting.

The action a_t^i will lead to stochastic transition from the current state to the next, in part because the taxis receiving our notifications will decide whether to accept or ignore our request, depending on the zone they are currently in. These decisions will then lead to movement of vacant taxi (the supply). The variable otw_t^i captures the number of taxis on the move, while $W_{v_t}^j$ captures taxis entered and currently waiting at the queue. Finally, part of the state transition also brings in new passengers, which is denoted as $W_{p_t}^j$.

Based on the state space definition above, and the uncertainties involved, we define the transition probability as:

$$T(s_{t+1}, a_t, s_t) = \prod_i Prob_{accept}^i(Accept_t^i) \cdot Prob_{demand_t}.$$

As will be discussed later, we assume that the demand arrival follows a Poisson arrival process with known and time-dependent arrival rates.

Finally, the reward function is defined as:

$$\begin{aligned} R(s_{t+1}, s_t, a_t) = \quad & -\textstyle\sum_{j=1}^{5}(\gamma_v \cdot W_{v_{t+1}}^j + \gamma_p \cdot W_{p_{t+1}}^j) \\ & -(\lambda_v \cdot W_{v_{t+1}}^6 + \lambda_p \cdot W_{p_{t+1}}^6). \end{aligned} \tag{1}$$

TABLE IV

NOTATIONS

Notations	Description
v	Vehicles.
p	Passengers.
t	Time period.
s_t	state at epoch t.
a_t	Action taken at epoch t.
i	zone representation.
gap^{iter}	Demand and supply gap at iteration = iter.
gap_c^{iter}	Cumulative Demand and supply gap at iteration = iter.
$iter$	Iterations for demand input.
$Dist_t^i$	Number of available taxis in zone i at time t.
a_t^i	Number of taxis we are going to engaged into airport trips at time t for zone i.
otw_t^i	Number of taxis moving towards airport at time t and zone i.
$supply_t^i$	Number of taxis that haven't been notified in zone i at time t.
${W_p}_t^j$,$j = 0, 1, ..., 6$	At passenger queue, number of passengers waited for j periods at time t.
${W_v}_t^j$: $j = 0, 1, ..., 6$	At taxi queue, number of taxis waited for j periods at time t.
$Accept_t^i$	Number of accepted taxis in zone i at time t.
$Prob_{accept}^i$	Probability of acceptance in zone i.
$Prob_{demand_t}$	Probability of demand inflow at t.
$T(s_{t+1}, a_t, s_t)$	Transition probability of moving from state s_t to state s_{t+1}, after taking action a_t.

The γ_v and γ_p values are the penalty for waiting one period, while the λ_v and λ_p values are the penalty for excessively long waiting time (> 6 periods). We assume that $\lambda > \gamma$. In our numerical study we assume that $\gamma_v = \gamma_p$ and $\lambda_v = \lambda_p$ for simplicity. But it can be easily tweaked by the planner if certain emphasis needs to be placed on either passengers or taxi drivers.

B. Dual LP Formulation

To actually solve the MDP, we first convert the MDP formulation into the following dual linear programming formulation, and solve it using the commercial optimization software such as CPLEX.

$$\begin{aligned} \max_x \quad & \sum_{t,s,a} R^t(s,a)x^t(s,a) \\ \text{s.t.} \quad & \sum_a x^0(s',a) = \delta(s'), \forall s' \\ & \sum_a x^{t+1}(s',a) - \sum_{s,a} x^t(s,a)T^t(s,a,s') = 0, \\ & \qquad \forall t \geq 0, \forall s' \\ & x^t(s,a) \geq 0 \end{aligned} \tag{2}$$

After solving the LP dual, the policy is obtained by normalizing $x^t(s,a)$: $\pi^t(s,a) = \frac{x^t(s,a)}{\sum_{a'} x^t(s,a')}, \forall t, s$. Here x represents number of times actions are executed in the states.

V. NUMERICAL EXPERIMENTS

The most important point we would like to demonstrate in our numerical experiments is the superiority of integrating both the demand prediction and the control optimization on the taxi supply. To show this, we conduct the following three sets of experiments, all of them using the same randomly generated demand instances to ensure that their results are directly comparable.

- Optimal control with predicted demand.
- Optimal control with Poisson demand assumption.
- No control.

By comparing the first two sets of experiments, we can observe the effectiveness of integrating demand prediction into the optimal control by solving MDP. With the third set of experiments, we can observe how significant it would be to have controls of different levels of sophistication.

In the first two experiment configurations, we use the real-world datasets collected from October and November of 2019 for training demand models. For all three configurations, we then use the data from December of 2019 for the testing purpose. To quantify the solution quality, we calculate the cumulative absolute gap between passenger demands and taxi supplies over the whole experimental horizon.

A. Experiment Setup

The size of the state space grows rapidly in response to the number of zones and the number of taxis under control. To optimally solve the MDP without having to rely on heuristics, we define our problem setup to contain two zones, with two time periods. Zone 1 is defined to contain all areas that are within 15-minute driving time to the airport, while Zone 2 contain all outer regions of interest that have more than 15-minute driving time. The transition probabilities between zones and the airport queues are derived from the real-world dataset.

We assume that $Prob_{accept} = 0.7, 0.6$, and $\lambda = 1$ and $\gamma = 3$ for both demands and supplies in the reward function. We set the number of taxis in Zone 2 to be 100, while varying the number of taxis in Zone 1 to be: 10, 30, and 50.

The experiment results are summarized in Table V. In the *gap* column, we summarize the sum of absolute demand/supply gaps over the planning horizon. In all problem sizes, we can clearly see that **optimal control with demand prediction** outperforms both **optimal control with Poisson assumption** and **No control** by a large margin. And a surprising finding is that even with optimal control, if we adopt not-so-ideal demand model (Poisson arrival in this case), the control policy we obtain can lead to inferior results than no control.

By looking at the largest numerical case (50,100), we can see that our proposed approach of optimal control with demand predictions outperforms the same control strategy, yet with Poisson demand assumption, by 43%. Against the status quo, which has no external control, we could reduce the gap by 23%.

To further understand where does the gap comes from, we also breakdown the measured absolute gap into *demand gap* and *supply gap*, to indicate whether we have too many unserved demands (indicated by 'gap (demand)') or too many idle supplies (indicated by 'gap (supply)').

TABLE V

ONE ZONE AND OUTER ZONE WITH SIMULATED ACCEPTANCE (T=2)

type	distribution	gap	gap (demand)	gap (supply)
Optimal control w/ predicted demand	(10,100)	13.15	92.86%	7.14%
	(30,100)	12.25	82.48%	17.52%
	(50,100)	11.29	73.99%	26.01%
Optimal control w/ Poisson assumption	(10,100)	22.65	94.86%	5.14%
	(30,100)	20.31	91.49%	8.51%
	(50,100)	19.80	86.29%	13.71%
No control	(10,100)	16.17	95.66%	4.34%
	(30,100)	15.15	89.04%	10.96%
	(50,100)	14.62	83.41%	16.59%

VI. CONCLUSION

In this paper, we study the problem of how to optimally control taxi queues at demand hotspots such as airports. By using one of the busiest airports in the world as an example, we demonstrate the importance of integrating both demand predictions and optimal control in balancing passenger demands and taxi supplies at the airport taxi queues. Our contributions in this paper are two-fold: on one hand, we propose a simple yet effective way in predicting passenger demands, which rely on the use of cumulative flight arrivals, and a number of flight-related information. The chosen features are so powerful that a simple linear regression model is sufficient for us to have accurate demand predictions. In the optimal control section, we demonstrate that it is critically important to have good demand prediction information, since the solution to the MDP-powered optimal control policy will be severely affected if inaccurate demand information is incorporated. In our numerical studies based on real-world data, we observe that our proposed approach of optimal control with demand predictions outperforms the same control strategy, yet with Poisson demand assumption, by 43%. Against the status quo, which has no external control, we could reduce the gap by 23%. These results demonstrate that our proposed methodology has strong real-world potential.

Although major airports around the globe are unfortunately being affected greatly by the COVID-19 pandemic, we believe that our model would be particularly relevant when the traffic gradually restores at major airports.

REFERENCES

[1] H. Yao, F. Wu, J. Ke, X. Tang, Y. Jia, S. Lu, P. Gong, J. Ye, and Z. Li, "Deep multi-view spatial-temporal network for taxi demand prediction," in *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.

[2] X. Geng, Y. Li, L. Wang, L. Zhang, Q. Yang, J. Ye, and Y. Liu, "Spatiotemporal multi-graph convolution network for ride-hailing demand forecasting," in *Thirty-Third AAAI Conference on Artificial Intelligence*, 2019.

[3] Y. Lu, S. Xiang, and W. Wu, "Taxi queue, passenger queue or no queue?" in *18th International Conference on Extending Database Technology*, 2015, pp. 593–604.

[4] M. S. Rahaman, M. Hamilton, and F. D. Salim, "Queue context prediction using taxi driver knowledge," in *Proceedings of the Knowledge Capture Conference*, 2017, pp. 1–4.

[5] M. S. Rahaman, Y. Ren, M. Hamilton, and F. D. Salim, "Wait time prediction for airport taxis using weighted nearest neighbor regression," *IEEE Access*, vol. 6, pp. 74 660–74 672, 2018.

[6] G. L. Curry, A. De Vany, and R. M. Feldman, "A queueing model of airport passenger departures by taxi: Competition with a public transportation mode," *Transportation Research*, vol. 12, no. 2, pp. 115–120, 1978.

[7] S.-F. Cheng and X. Qu, "A service choice model for optimizing taxi service delivery," in *Twelfth International IEEE Conference on Intelligent Transportation Systems*, 2009, pp. 1–6.

[8] C. Kamga, A. Conway, A. Singhal, and A. Yazici, "Using advanced technologies to manage airport taxicab operations," *Journal of Urban Technology*, vol. 19, no. 4, pp. 23–43, 2012.

[9] A. Anwar, M. Volkov, and D. Rus, "Changinow: A mobile application for efficient taxi allocation at airports," in *16th International IEEE Conference on Intelligent Transportation Systems (ITSC 2013)*. IEEE, 2013, pp. 694–701.

[10] S. S. Jha, S.-F. Cheng, M. Lowalekar, W. H. Wong, R. R. Rajendram, T. K. Tran, P. Varakantham, N. Truong Trong, and F. Abd Rahman, "Upping the game of taxi driving in the age of Uber," in *Thirtieth Annual Conference on Innovative Applications of Artificial Intelligence*, 2018, pp. 7779–7785.

[11] S.-F. Cheng, S. S. Jha, and R. Rajendram, "Taxis strike back: A field trial of the driver guidance system," in *International Conference on Autonomous Agents and MultiAgent Systems*, 2018, pp. 577–584.