Seeking in Ride-on-Demand Service: A Reinforcement Learning Model With Dynamic Price Prediction

Suiming Guo[®], Baoying Deng[®], Chao Chen[®], Jintao Ke, Jingyuan Wang[®], Saiqin Long, and Ke Xu[®], *Fellow, IEEE*

Abstract-Recent years witness the increasing popularity of Ride-on-Demand (RoD) services, such as Uber and Didi. Compared with traditional taxi, RoD service is more "datadriven" and adopts dynamic pricing to manipulate the supply and demand in real time. Dynamic price could be viewed as an accurate and quantitative indicator of the supply and demand, and could provide clues to drivers, passengers, and the service providers, possibly reshaping the ways in which some problems are solved. In this article, we focus on the seeking route recommendation problem that aims at increasing driver revenue by recommending highly profitable seeking routes to drivers of vacant cars with the help of dynamic prices. We first justify our motivation by showing the importance of route recommendation and answering why it is necessary to consider dynamic prices, based on the analysis of real service data. We then design a dynamic price prediction model to generate the dynamic prices at any given time and location based on multisource urban data. After that, a reinforcement learning model is adopted to perform seeking route recommendation based on predicted dynamic prices. We conduct extensive experiments in different spatiotemporal combinations and make comparisons with multiple baselines. Results first show that our dynamic price prediction model achieves an accuracy ranging from 83.82% to 90.67% under different settings. It also proves that considering the real-time predicted dynamic prices significantly increases driver revenue by, for example, 12% and 47.5% during weekday

Manuscript received 2 April 2024; revised 29 April 2024 and 26 May 2024; accepted 27 May 2024. Date of publication 30 May 2024; date of current version 6 September 2024. This work was supported in part by the Guangdong Basic and Applied Basic Research Foundation under Grant 2024A1515012094; in part by the National Natural Science Foundation of China under Grant 62002135, Grant 62322601, Grant 62425201, and Grant 62172066; in part by the Excellent Youth Foundation of Chongqing under Grant CSTB2023NSCQJQX0025; and in part by the National Natural Science Foundation of China under Grant 72171013. (*Corresponding author: Chao Chen.*)

Suiming Guo, Baoying Deng, and Saiqin Long are with the National and Local Joint Engineering Research Center of Network Security Detection and Protection Technology, Guangdong Provincial Key Laboratory of Data Security and Privacy Protection, and the College of Information Science and Technology, Jinan University, Guangzhou 510632, China (e-mail: guosuiming@email.jnu.edu.cn).

Chao Chen is with the College of Computer Science, Chongqing University, Chongqing 400044, China (e-mail: cschaochen@cqu.edu.cn).

Jintao Ke is with the Department of Civil Engineering, University of Hong Kong, Hong Kong, SAR, China (e-mail: kejintao@hku.hk).

Jingyuan Wang is with the School of Computer Science and Engineering and the School of Economics and Management, Beihang University, Beijing 100191, China (e-mail: jywang@buaa.edu.cn).

Ke Xu is with the Department of Computer Science and Technology, Tsinghua University, Beijing 100084, China (e-mail: xuke@mail.tsinghua.edu.cn).

Digital Object Identifier 10.1109/JIOT.2024.3407119

evening rush hours, than merely using the average prices or completely ignoring dynamic prices.

Index Terms—Driver revenue, dynamic price, reinforcement learning, Ride-on-Demand (RoD).

I. INTRODUCTION

D URING the last decade, Ride-on-Demand (RoD) services, such as Uber and Didi, first emerged as a new and disruptive mode of transport compared with traditional taxi service and have gained increasing popularity since then. RoD service is beneficial for passengers and drivers. For passengers, it is more convenient to request for service through smart phones, and the price is more affordable; for drivers, in many cities, it is not required to apply for licenses or medallions before driving, and working schedules are more flexible.

RoD service has two unique features, namely, the use of *dynamic pricing*, and being *data-driven*.

Dynamic pricing, also known as "surge pricing" when it was introduced in Uber, aims to manipulate the supply and demand in real time, i.e., the service provider sets a higher price when demand—potential passenger requests—exceeds supply—vacant cars on the road, and vice versa. Possible forms of dynamic pricing include auction-based mechanisms, additive bonus to the trip fare, and multiplying the trip fare by a real-time factor (called "dynamic price multiplier"), etc. Most service providers choose the multiplicative form in practice, and so does our study.

Dynamic pricing could be viewed as a closed-loop feedback mechanism to manipulate the supply and demand: the price is determined based on the real-time supply and demand condition, and it, in turn, controls the supply and demand in a way leading to a higher service efficiency. Besides, dynamic prices could also be regarded as an accurate indicator of supply and demand.

Data-Driven: The use of on-board GPS devices triggers a large amount of studies and applications in taxi service since the beginning of the last decade. Examples of studies based on GPS trajectories include inferring order origins and destinations, detecting possible detours, optimizing for driver revenue based on driving habits, etc.

RoD service goes one step further and tremendously expands the volume and variety of data. It is mainly based on smart phones and mobile apps, which is a data source providing much more abundant information. For example,

2327-4662 © 2024 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

accurate order information is now readily available, including origin, destination, the time getting on/off the car, the time and location the request is issued, text descriptions (typed by the user on the mobile app) of origin and destination, etc. User behavior log is also possible, as it could be recorded by the mobile app. Typical user behaviors include the type of mobile device, the number of times one checks for trip fare, the regions one zooms in and out to view available drivers, etc. All these new sources of data offer richer information than the traditional GPS trajectories.

With these two new features, it is now possible to reshape the ways in which some traditional problems are solved. In this article, we focus on the seeking route recommendation problem, a widely studied problem in mobility service, especially taxi. Basically, seeking route recommendation aims at increasing driver revenue by recommending profitable seeking routes to drivers of vacant cars. Drivers themselves may adopt some naive strategies based on experience or word-of-mouth. Such strategies may be inaccurate, and may even suggest a large number of drivers blindly flocking to the same hot spots such as the central business district, exacerbating the already imbalanced supply and demand condition. Algorithms developed in previous taxi-related studies solve this problem by finding local or global hot spots, modeling driver behavior using Markov decision process (MDP), simulating driver behavior using physics-based approaches, etc.

Things are different in RoD service. First, the data-driven feature enables us to acquire more data sets, such as multisource urban data, to solve the problem. Second, dynamic prices could be an accurate and readily available indicator of supply and demand, which previous studies go into great lengths to find out. Furthermore, to make full use of dynamic prices that fluctuate in real time, it should be evaluated on a fine spatiotemporal granularity, e.g., using recent or average dynamic price multipliers may not be enough.

In this article, we solve the seeking route recommendation problem by combining reinforcement learning with dynamic price prediction. We first give intuitive answers to two questions, i.e., why recommending seeking routes and why considering dynamic prices, based on the analysis of real service data. After that, a seeking route recommendation framework consisting of two parts is built. The first part is a dynamic price prediction model based on multisource urban data, trying to obtain a price multiplier given the spatiotemporal condition and other relevant features when a driver picks up a passenger. Evaluation results show that our dynamic price prediction model achieves an accuracy ranging from about 83% to 90% in different settings, laying a solid foundation for the reinforcement learning model. In the second part, we use MDP to model the seeking behavior, and adopt a reinforcement learning model to tackle the seeking route recommendation problem, in which the predicted dynamic prices from the first part are incorporated into reward design. We evaluate driver revenue and utilization rates under our model, and it proves that taking the real-time predicted dynamic prices into consideration significantly increases driver revenue than merely using average dynamic prices or completely ignoring dynamic prices.

Our contributions are listed in the following.

- We summarize the two unique features of RoD service, i.e., dynamic pricing and data-driven, and solve a common problem, albeit widely studied in previous taxi service, by making full use of these two features. By comparison, most previous works do not consider dynamic pricing and simply treat RoD service as a new version of taxi. Even the few works that indeed consider dynamic pricing only calculate the average or historical prices and ignore its real-time nature.
- 2) We adopt a reinforcement learning model to solve the seeking route recommendation problem, and incorporate dynamic prices. The reinforcement learning model helps to consider the long-term effects of seeking routes and thus redistributes drivers more effectively.
- We conduct extensive experiments and comprehensively evaluate both our dynamic price prediction and reinforcement learning model on real service data set.

The remainder of this article is organized as follows. Section II reviews the related work. In Section III, we present our data sets and provide a detailed and intuitive analysis. Sections IV and V discuss the dynamic price prediction and the reinforcement learning model, respectively. Section VI presents a comprehensive evaluation on both models. Section VII gives a brief summary and some discussions based on evaluation results. Finally, Section VIII concludes this article.

II. RELATED WORK

We provide discussions on related work about three topics: 1) RoD service; 2) dynamic price and its prediction in RoD service; and 3) seeking in taxi or RoD service.

RoD Service: RoD service is also known as on-demand ridehailing, and in some cases, people may call it as ride-sharing. In fact, "RoD" and "ride-sharing" have different emphases. Ride-sharing emphasizes the share of rides with the same or similar origins or destinations, either between passengers or between the driver and passengers. RoD, on the other hand, emphasizes "on-demand"—the service is available as soon as one asks for it at any time or location. Under this setting, RoD service is viewed as a disruptive new version of taxi, by supporting mobile app, adopting dynamic prices, and being data-driven. Our study focuses on RoD service, so we omit the discussions on ride-sharing.

Compared with taxi, RoD service is relatively new, and thus receives limited attention. Most studies simply treat RoD as a service similar to taxi, and want to find out the differences between them. For example, [1] focuses on the passenger waiting time and make a comparison, of both waiting time and price, between Uber and taxi; [2] and [3] choose the market share as the study target; and [4] and [5] discuss the impacts and market effects of Uber's entrance—e.g., how driver behavior is changed since Uber takes place.

Dynamic Price and Its Prediction in RoD Service: Dynamic pricing plays an important role in many services, as an effort to either improve service efficiency or manipulate supply and demand in different forms. For examples, dynamic pricing As a unique feature of RoD service, dynamic pricing and its effects are studied in a number of works. For example, [11], [12], and [13] discuss the effects of dynamic pricing in balancing and redistributing supply and demand, increasing driver revenue, and reducing passenger waiting time. Reference [14], as a typical early study of on-demand ride-hailing service, evaluates Uber's surge pricing as a blackbox by placing simulated mobile app users across important locations. References [15] and [16] explore demand pattern, the effects of dynamic pricing on passengers, and passenger behavior. Reference [17] combines pricing with dispatching, which is a more traditional problem, and proposes a distributed pricing framework. Some studies take an economics perspective: examples that consider the effects of dynamic pricing on supply and demand include [11], [18], and [19].

The problem of dynamic price prediction is also tackled using various methods. For example, [20] defines the predictability of price multipliers and uses Markov-chain or neural network models to predict average dynamic price multiplier of a region based on the predictability. References [21] and [22] turn to the prediction at a finer granularity and predict the dynamic price multiplier given time and location using linear regression or neural network based on multisource urban data. Reference [23] summarizes the above works, adopts an ensemble learning model, and chooses different models based on price multiplier predictability. Reference [24] takes a different path, emphasizes the interpretability of price prediction results, and presents a simple but quantifiable approach to dynamic price prediction. It gives a detailed evaluation as to what features have more obvious impacts on price multipliers.

Seeking in Taxi or RoD Service: Seeking analysis is also a highly visible topic in traditional taxi service, and in RoD service, it also receives some, yet limited, attention. Generally speaking, we could divide such analysis into seeking strategies analysis and seeking route recommendation, the latter of which is the target of study in this article.

Seeking strategies analysis could be considered as a macrolevel problem, and tries to uncover the relationship between driver seeking strategies (e.g., choosing hot spots, driving faster, etc.) and revenue. For example, [25] and [26] consider two different strategies, i.e., hunting or waiting for passengers, and compare their performances under different circumstances, based on taxi GPS trajectories. In RoD service, [27] collects multisource urban data and designs a framework to mine the relationship between driver revenue and the carefully crafted features that are relevant to seeking strategies. In such relationship, dynamic prices become an important component.

Seeking route recommendation could, on the other hand, be treated as a micro-level problem, and aims to recommend the right road segment or city cell a driver should keep seeking for so that driver revenue is increased. In these studies, MDP is frequently used to model the interaction between drivers and the service itself—e.g., [28], [29], [30], and [31]. Another example using MDP is [32], and it mainly pays attention to electric taxis and make charging decisions based on both battery constraint and GPS trajectories. Besides MDP, [33]

generates recommendation results by minimizing the distance between taxis and potential passenger requests; [34] and [35] apply reinforcement learning; [36] solves the problem by allowing a single driver to be matched to multiple passengers; and [37] builds theoretical models and optimization problems. In RoD service, [38] uses *Q*-learning to recommend profitable seeking routes.

Different from the above works, our study combines the power of dynamic pricing and the data-driven feature to conduct seeking route recommendation. Our methodologies are justified by analyzing real service data. The dynamic price prediction model is built on multisource urban data and achieves a satisfactory prediction accuracy. The prediction results help to provide references to drivers of vacant cars, based on which the reinforcement learning model generates seeking route recommendations. Extensive experiments are carried out to evaluate the effects of using a reinforcement learning model and of considering predicted dynamic prices.

III. DATA AND ANALYSIS

We first present the multisource urban data sets used in this article, and then provide a detailed data analysis to give intuitive answers that motivate our study.

A. Multisource Urban Data Sets

1) RoD Data Sets: Unlike previous studies on seeking route recommendation in taxi service that mainly rely on taxi GPS trajectories, in RoD service, the "data-driven" feature makes it possible to obtain more data sets. In our study, we also acquire order data and dynamic price data besides GPS trajectories.

GPS Trajectories: This is similar to the widely used taxi GPS trajectories data set. It contains the periodic GPS records, in longitude and latitude, of every single car in operation. Each record includes the longitude and latitude, time stamp, speed, direction, the unique car ID, etc. This data set spans from August to November 2016, and contains the records of about 3500–3800 cars daily in Beijing, China. Specifically, the ranges of longitude and latitude are [116.21, 116.56] and [39.81, 40.08].

Order and Dynamic Price Data: The information of orders and dynamic prices are recorded in the same data set. For the order data, in RoD service, the use of mobile app enables accurate recording of order information, as order origins and destinations are now clearly specified by users and recorded by the service provider. For the dynamic price data, the price multiplier is recorded when the service provider returns the estimated trip fare and current price multiplier. Sometimes getting the trip fare estimates does not mean order creation, and to accurately associate an order with a price multiplier, only the price multiplier returned in the trip fare estimate that is closest to order creation is retained. This data set also covers same time span, and the total number of orders is 2742120. Each entry includes origin, destination, the time getting on and off, the unique ID of passenger/driver/car/order, the estimated trip fare, price multiplier, etc.



Fig. 1. Probability distribution functions of revenue efficiency during four representative time periods. (a) [8:00, 9:00). (b) [12:00, 13:00). (c) [18:00, 19:00). (d) [23:00, 0:00).



Fig. 2. Probability distribution functions of searching time during four representative time periods. (a) [8:00, 9:00). (b) [12:00, 13:00). (c) [18:00, 19:00). (d) [23:00, 0:00).

In all the above data sets, all unique IDs of drivers, passengers, and cars are anonymized so that one cannot relate an ID to a specific person or car. All the above data sets are obtained from Shenzhou UCar, a major RoD service provider in China.

2) Taxi Data Set: Besides RoD data sets, we also use taxi GPS trajectories as an auxiliary data set. The reasons are straightforward. First, as mentioned in [24], RoD and taxi service are similar and complementary to each other, and hence driver behavior should be similar. For example, a hot region in taxi service may also be a good choice for RoD drivers. In other words, taxi GPS trajectories could serve as a useful guideline for RoD drivers. Second, taxi service data also describes car movements on the roads and is helpful in characterizing the traffic in different spatiotemporal combinations. For example, if the average speed of taxis is low, then it may be an indication of traffic congestion in a region. Our taxi GPS trajectories covers the same time span and contains records of about 30 000 taxis in Beijing.

3) Public Data Sets: Public data sets are used to provide more information. There are a lot of possible choices, and in our study, we choose the POI data and public transport distribution data, as described in the following.

POI Data: POI information is widely used in studies on location-based service, and it characterizes the type of a location, such as airport, restaurant, etc. We rely on POI data to describe the origins and destinations of orders. Our data is crawled from AMap service [39], one of the most popular online map service providers in China. Each POI falls in one of the 14 categories: car service, restaurant, shopping, sports and entertainment, hospital, hotel, scenic spot, business and residential building, government, education and culture, transportation facility, finance and insurance, business, and lifestyle.

Public Transport Distribution Data: Public transport services—e.g., bus, metro, tram, and public rental bike—also

play an important role in transportation, and the effects of considering public transport data are similar to that of considering taxi data. That is, the status of public transport services could give a hint on RoD driver behavior and describe traffic condition. We choose the distribution of bus and metro service as our public transport distribution data. Specifically, we acquire the locations of all bus (and metro) stations and all the buses or metros that stop by these stations, from AMap service.

B. Data Analysis

We give intuitive answers to two questions, i.e., *why recommending seeking routes* and *why considering dynamic prices*. They not only help to form a comprehensive picture of RoD service but also provide inspiration for our study.

1) Why Recommending Seeking Routes?: We first give the definition of driver revenue. The fare of an RoD trip, i.e., the driver revenue during this trip, is defined as

$$r = dp_o \cdot (f_{\text{base}} + f_d * d_{\text{trip}}). \tag{1}$$

In (1), dp_o , f_{base} , and f_d are the dynamic price multiplier at the trip origin, the flag-fall price, and the unit price per kilometer, respectively. d_{trip} is the trip distance from the origin to destination. This is an approximation, ignoring the extra fare of slow driving and minimum distance, and we consider it acceptable because these two terms are usually much smaller. In our data set, we have $f_{\text{base}} = 15$ and $f_d = 2.8$, both in RMB Yuan.

We then show the distribution of revenue efficiency (i.e., the revenue obtained in unit time) among all drivers during four representative time periods, i.e., [8:00, 9:00), [12:00, 13:00), [18:00, 19:00), and [23:00, 0:00), in Fig. 1. Intuitively, these four time periods correspond to morning rush hour, nonrush hour around noon, evening rush hour, and late night hour. Similarly, in Fig. 2, we also show the distribution of searching time, i.e., the time a driver seeks for passengers between

 TABLE I

 Statistics of Revenue Efficiency and Searching Time During Four Representative Time Periods

Time period	Revenue	efficiency (Yua	an/min)	Searching time (min)			
	1 st quartile	3 rd quartile	average	1 st quartile	3 rd quartile	average	
[8:00, 9:00)	0.71	1.06	0.88	8	20	15.49	
[12:00, 13:00)	0.62	0.96	0.80	15	29	22.63	
[18:00, 19:00)	0.63	0.97	0.81	12	26	19.73	
[23:00, 0:00)	0.54	0.98	0.78	20	36	28.14	

two consecutive trips, during these time periods. In addition, Table I summarizes some simple statistics—i.e., the 1st quartile, 3rd quartile, and average—of both revenue efficiency and searching time from Figs. 1 and 2.

We have the following observations.

- Both revenue efficiency and searching time vary significantly across different time periods.
 - a) For example, in the morning rush hour, the average revenue efficiency and average searching time are the highest and the lowest, respectively, among all four time periods. The late night hour is just the opposite.
 - b) Comparing between the late night hour and the evening rush hour (or the nonrush hour around noon), it should also be noted that the difference of searching time is much more obvious than that of revenue efficiency. This makes the seeking route recommendation problem more complicated.
- 2) Both revenue efficiency and searching time vary significantly across drivers. For revenue efficiency, during [8:00, 9:00), the 3rd quartile is 49.30% higher (1.06 versus 0.71) than the 1st quartile, and during [23:00, 0:00), the percentage is much higher and achieves 81.48% (0.98 versus 0.54). For searching time, the comparison is similar.
- 3) Therefore, it is necessary to recommend seeking routes. First, a good seeking strategy or seeking route could help those lower earning drivers become more efficient and earn more. Second, such recommendation should consider carefully as to what features influence the recommendation results. For example, the above observations show that the temporal features are important.

2) Why Considering Dynamic Prices?: Previous studies on seeking route recommendation in taxi service usually go to great lengths in finding "hot spots" or finding locations with more high-earning orders. In RoD service, the dynamic price multiplier is an accurate indicator of the supply and demand condition—which is just the goal of using dynamic pricing—and it should be an integrated description of the information that previous studies want to find out.

To justify that it is necessary to consider dynamic prices in seeking route recommendation, in the following, we show and explain the distribution of dynamic price multipliers on both temporal and spatial dimension. Fig. 3 first shows the temporal distribution of dynamic price multipliers on both weekdays and weekends. In Fig. 3, we calculate the average dynamic price multiplier among all orders every half an hour, across the whole city of Beijing. It is clear that as follows.



Fig. 3. Temporal distribution of dynamic price multipliers.

- The average dynamic price multiplier is greater than 1.0 in most of the time, and thus dynamic prices have a significant impact on driver revenue. In some time periods, e.g., the evening rush hours on weekdays, the average dynamic price multiplier reaches as high as 1.35.
- 2) The dynamic price multiplier has different patterns during different hours-of-day or days-of-week, and hence temporal features are important. For example, on weekdays, there are four obvious peaks: the small hours, the morning and evening rush hours, and the late night hours. On weekends, the patterns are different: there are not morning rush hours, and price multipliers remain at a low level until evening rush hours.
- 3) The dynamic price multiplier fluctuates rapidly throughout a day. This is the result of its real-time nature, as it is designed to reflect the real time changes of supply and demand. As a result, when considering dynamic prices in our study, a fine temporal granularity is needed, and simply calculating the average or historical price multipliers may not be enough.

On the spatial dimension, we focus on the evening rush hours [18:00, 19:00) on weekdays, divide the area of study (i.e., within longitude [116.21, 116.56] and latitude [39.81,40.08]) into 900 (= 30×30) cells, and plot the spatial distribution of the number of orders (i.e., *demand*), dynamic price multipliers, and the number of vacant cars (i.e., *supply*) in Figs. 4–6. It is shown that as follows.

- 1) The distributions of the number of orders as well as vacant cars are spatially imbalanced, and it is thus necessary to recommend seeking routes. This agrees with common intuitions: such numbers are much higher in the city center, and are drastically reduced in city suburbs.
- 2) The distribution of dynamic price multipliers is also spatially imbalanced, but with more complicated patterns. Intuitively, the price multiplier is higher in the



Fig. 4. Spatial distribution of the number of orders.



Fig. 5. Spatial distribution of dynamic price multipliers.



Fig. 6. Spatial distribution of the number of vacant cars.

city center. Counter-intuitively, the price multiplier in the city suburbs surrounding the city center is even higher, indicating severe supply and demand imbalance.

- 3) Comparing the above two observations, the reason of having higher price multipliers but fewer orders and cars in city suburbs is that drivers tend to flock to city center to seek for passengers, leaving a very limited number of drivers in suburbs. Though the orders are indeed fewer in suburbs, the number of drivers is still not enough to meet the demand. To the contrary, for those drivers going to city center, though there are indeed more orders, there are even more drivers, making price multipliers lower.
- 4) Therefore, if dynamic price multipliers are observed and considered in seeking route recommendation, it is possible to guide more drivers to make better and more informed decisions, e.g., staying in city suburbs to seek for passengers. This not only increases driver revenue,

but eases the imbalance between supply and demand and improves service efficiency as well.

To summarize, the reasons why dynamic prices should be considered in seeking route recommendation are twofold.

- Dynamic price multiplier changes rapidly on the temporal dimension and has a significant impact on driver revenue, so it should be carefully considered in a fine granularity.
- The spatial distribution of dynamic price multipliers reveals some counter-intuitive insights, showing that considering dynamic prices is helpful in guiding drivers to make better and more informed decisions.

IV. DYNAMIC PRICE PREDICTION

The rationale behind doing dynamic price prediction is simple. First, dynamic price multiplier should be considered in seeking route recommendation so as to guide drivers to make better decisions. Second, dynamic price multiplier should be evaluated at a fine spatiotemporal granularity, so simply using average or historical price multipliers may not be enough.

In RoD service, the price multiplier is usually discrete. For example, in our data set, the price multiplier falls in the range $DP_{range} = [1.0, 1.1, 1.2, 1.3, 1.4, 1.5, 1.6]$. Therefore, predicting the price multiplier could be viewed as a classification problem. The input to the classifier is a feature vector extracted from our multisource urban data that gives further details based on the spatiotemporal information. The output is one of the 7 possible price multipliers. We have the following definition of our dynamic price prediction problem.

Definition 1 (Dynamic Price Prediction Problem): Given the spatiotemporal information (e.g., time, date, longitude, latitude, etc.).

Extract an input feature vector, denoted by \tilde{X} , from our RoD order and dynamic price data, taxi data, POI data, and public transport distribution data.

Predict $\hat{p}(y = \bar{d}|\bar{X}) \quad \forall \bar{d} \in DP_{\text{range}}$: the probability of a candidate price multiplier \bar{d} being the actual price multiplier *y* based on the input feature vector \vec{X} . The price multiplier with the largest probability is the output of the classification problem.

A. Feature Extraction

In feature extraction, we first obtain the spatiotemporal information (e.g., time, date, longitude, and latitude) from every order in our RoD order data, and then generate the following features corresponding to each order. Features are divided into four different contexts, namely, temporal context, spatial context, historical dynamic prices context, and taxi service context. The price multiplier obtained from our RoD dynamic price data that corresponds to each order is used as the ground-truth in model training.

1) Temporal Context: Temporal context features are the most basic and are simply the time and date obtained from each order in RoD order and dynamic price data.

Time Feature ET: The hour and minute value (i.e., h and min) are mapped to the radian of a unit circle. On the circle,

the 1440 min of a day are represented by a radian value between $[0, 2\pi)$. The time feature *ET* is defined as

$$ET = [\cos\theta, \sin\theta], \theta = \frac{h * 60 + \min}{1440} * 2\pi.$$
(2)

Date Features DM, DW, and DH: DM, DW, and DH refer to the day of month, the day of week, and a Boolean value describing if the day is a weekday, respectively.

2) Spatial Context: Features describing the spatial context are extracted mainly from the POI data and public transport distribution data. Based on the longitude and latitude, we describe a location by the POIs and the availability of public transport around it.

POI Features POI_n, POI_f, and POI_u: They are all 14-D vectors, and reflect the number, frequency and uniqueness of POIs around a location. The 14 dimensions correspond to the 14 POI categories explained in Section III-A3. We count the POIs within a 500-m radius of the location. For the *i*th category of POI, the number of POIs of this category around the location is denoted by POI_{n,i}, and the total number of POIs of this category in the city is M_i , then the frequency and uniqueness of the *i*th category of POI—i.e., POI_{f,i} and POI_{u,i}—are defined as

$$\operatorname{POI}_{f,i} = \frac{\operatorname{POI}_{n,i}}{\sum_{k=1}^{14} \operatorname{POI}_{n,k}}$$
(3)

$$\mathrm{POI}_{u,i} = \frac{\mathrm{POI}_{n,i}}{M_i}.$$
(4)

Among these three vectors, POI_n simply counts the number of POIs of different categories around the location; POI_f focuses on the proportion of each category of POI to all POIs around the location; and POI_u characterizes the proportion of the number of a particular category of POIs to the total number of this category in the whole city, which is a reflection of POI uniqueness. Then, the POI feature vector is written as

$$POI = [POI_n, POI_f, POI_u].$$
(5)

Public Transport Distribution Feature BM: BM is a 4-D vector, describing the number of bus stations, the number of bus lines, the number of metro stations and the number of metro lines within a 500-m radius of the location. It is intuitive that the public transport distribution nearby has impacts on dynamic prices. First, it reflects the popularity of a location. Second, the presence of bus and metro stations also makes it possible for RoD drivers to provide connecting services—picking up a passenger who just alighted from a bus or train, or delivering a passenger to a bus or metro station—and hence changes passenger demand around a location.

3) Historical Dynamic Prices Context: Dynamic price multiplier is a perfect and real-time reflection, as claimed by major RoD service providers, of the supply and demand. As the current supply and demand may be related to the past supply and demand, we hypothesize that the past price multipliers should be helpful in predicting current price multipliers.

Based on the current time on which we want to evaluate the price multiplier, we extract the average price multiplier among all orders taking place within a 500-m radius of the location, during the last three half-hour-length timeslots

$$HDP = [APM_{-1}, APM_{-2}, APM_{-3}].$$
(6)

In (6), HDP is the historical dynamic prices feature vector, and $APM_{-i}(i = 1, 2, 3)$ are the average price multiplier during the last three timeslots, respectively. For example, if we want to predict the price multiplier at 17:09, then APM_{-1} , APM_{-2} , and APM_{-3} are the average price multiplier during [16:30, 17:00], [16:00, 16:30], and [15:30, 16:00], respectively.

4) Taxi Service Context: This context is extracted from taxi GPS trajectories. Each feature is calculated on a half-hourlength timeslot basis, and includes the taxis within a 500-m radius of the location. We extract the following features.

- 1) *Taxi Up Count UC:* The number of taxis that start new orders around the location during the current half-hour-length timeslot.
- Taxi Down Count DC: The number of taxis that terminate existing orders around the location during the current half-hour-length timeslot.
- 3) *Taxi Visit Count VC:* The number of taxis that present around the location during the current half-hour-length timeslot.
- 4) *Taxi Full Count FC:* The number of taxis with passengers on board that present around the location during the current half-hour-length timeslot.
- 5) *Taxi Full Ratio FR:* The ratio of taxi full count *FC* to taxi visit count *VC*.

Then, the taxi service feature vector is written as

$$TS = [UC, DC, VC, FC, FR].$$
(7)

5) Input Feature Vector: Before creating the input feature vector, it is necessary to normalize the above-mentioned features to guarantee convergence and a shorter training time. Among these features,

- 1) Date features *DM* and *DW* are categorical features, and they are represented by one-hot encoding. It is not necessary to do normalization.
- 2) Date feature *DH* is a Boolean, and takes the value of either 0 or 1. Hence, it is not necessary to normalize *DH*.
- 3) For all other features, we calculate the Z-score (i.e., the number of standard deviations from the mean) of each feature, or of each component of multidimension features, to perform normalization.

For each order, we gather all the above-mentioned features, as summarized in Table II, to generate the input feature vector \vec{X} , which are then fed into the dynamic price prediction model

$$X = [ET, DM, DW, DH, POI, BM, HDP, TS].$$
 (8)

B. Model Selection

There are a lot of available algorithms for a predictive task, including the popular deep learning models or more traditional and simpler machine learning models. Though deep learning models have achieved a tremendous progress on voice, text, or image data sets, they may not be good enough for our

Context	Feature	Dim.				
	ET: the time feature that maps the hour and minute to the radian of a unit circle					
Temporal Context	<i>DM</i> : one-hot encoded vector of the day-of-month value					
Temporal Context	DW: one-hot encoded vector of the the day-of-week value	7				
	DH: a boolean value that equals 1 when the day in question is a weekday	1				
	POI_n : the number of POIs, of 14 different categories, around the location	14				
Spatial Contaxt	POI_f : the POI frequency vector, as defined in (3), around the location					
Spatial Context	POI_u : the POI uniqueness vector, as defined in (4), around the location					
	BM: the number of bus/metro lines and stations around the location	4				
Historical	APM_{-1} : the average price multiplier around the location during the last timeslot	1				
Dynamic Prices	APM_{-2} : the average price multiplier around the location during the second last timeslot	1				
Context	APM_{-3} : the average price multiplier around the location during the third last timeslot	1				
	UC: taxi up count around the location during the current time slot	1				
	DC: taxi down count around the location during the current time slot	1				
Taxi Service	VC: taxi visit count around the location during the current time slot					
Context	FC: taxi full count around the location during the current time slot					
	FR: taxi full ratio around the location during the current time slot	1				

 TABLE II

 Summary of Features Used in Dynamic Price Prediction

dynamic price prediction task. As an example, we compare between deep models and tree-based models. First, recent discussions in [40] and [41] point out that tree-based models perform better than deep learning models on tabular data. Second, deep models are generally more complex and harder to fine tune, sometimes leading to overfitting problems. The more complex structure of deep models also gives rise to a longer training time. Finally, deep learning models are usually difficult to interpret or explain, and sometimes sophisticated techniques are needed. By comparison, tree-based models, especially those containing only a small number of trees, are naturally easier to interpret.

We thus choose tree-based models to conduct dynamic price prediction. Specifically, we adopt XGBoost and LightGBM models to predict the dynamic price multipliers, and perform a hit and trial to see which one leads to a higher prediction accuracy. In addition, we also implement a Random Forest model and an artificial neural network (ANN) model to serve as baselines, so that we could compare model performances and justify our model selection considerations. Below we briefly explain XGBoost, LightGBM, and Random Forest.

XGBoost: XGBoost [42] is a widely used gradient boosting decision tree variant that trains multiple base learners to improve model accuracy. It is highly parallelized and carefully optimized for a shorter training time. It is also currently one of the best, in terms of accuracy and training time, boosting-based tree-models.

LightGBM: LightGBM [43] is another widely used gradient boosting decision tree variant. But, compared with XGBoost, LightGBM focuses on efficiency and significantly reduces training time and memory consumption by adopting techniques, such as exclusive feature bundling (EFB), histogram-based algorithm, leaf-wise tree growth, gradientbased one-side sampling (GOSS), etc. On the other hand, LightGBM is not as robust as XGBoost is, and may be overfitting in some cases.

Random Forest: Random Forest is a bagging-based decision tree variant. Though it appears earlier than XGBoost and LightGBM, it is still widely used in many real world scenarios.

The bagging nature makes it especially suitable for parallel computation and scale well on large-scale high-dimensional data sets. Random Forest samples with replacement from the input data, and also uses feature bagging, or random subspace method, to ensure convergence and low correlation among multiple decision trees.

In Section VI, we would compare the performances of XGBoost, LightGBM, Random Forest, and ANN, in which the latter two are mainly used as baselines, and provide a detailed discussion on why certain models work better than others. Additionally, we also list the main parameters and their values used in these models.

V. REINFORCEMENT LEARNING MODEL

Now that we have built the dynamic price prediction model, the next task is to perform seeking route recommendation by a reinforcement learning model. Route recommendation could be conducted on different spatial granularities. For example, it could be done on the cell level, i.e., partitioning the area of study into rectangular cells and recommending the next cell a driver should go for seeking after the current cell, or on the road segment level, i.e., recommending the next road segment when a seeking driver comes to a road intersection. Route recommendation on the cell level may have a coarser spatial granularity, but has the advantage of being simple and efficiency, and still giving enough insights.

We choose to study the seeking route recommendation problem on the cell level. Similar to what we have done in data analysis in Section III-B2, we partition the area of study, i.e., the city of Beijing within longitude [116.21, 116.56] and latitude [39.81, 40.08], into 900 (= 30×30) rectangular cells of equal size. The time range is restricted to a whole day. In other words, we keep a timer τ and set it as $\tau = 0$ in the very beginning, and perform seeking route recommendation until $\tau \ge 1440$ (in minutes). The goal of the reinforcement learning model is to solve the following problem.

	TABLE III		
NOTATIONS USED IN OUR	REINFORCEMENT	LEARNING	Modei

X7 • • •	
Variable	Explanation
s, S	a state such that $s = [l, p, t, e]$, and S is the set
	of all states
l	the index of a cell
<i>p</i>	the index of current timeslot
P	the length of each timeslot (in minutes)
N	the total number of timeslots in a whole day
t	the number of minutes into the current timeslot
au	the timer that keeps the current time (in minutes)
e, E	e is the reverse of the incoming direction, and E
	is the set of all directions
a, A	a is the action a driver takes, and A is the set of
	all possible actions
$t_{seek}(j,p)$	the time used to seek for passengers in cell j
	during timeslot p
$d_{seek}(j,p)$	the driving distance to seek for passengers in cell
	j during timeslot p
$t_{drive}(j,k,p)$	the time used to drive from cell j to k during
	timeslot p
$d_{drive}(j,k,p)$	the driving distance from cell j to k during
	timeslot p
$P_{pickup}(j,p)$	the probability of picking up a passenger in cell
	j during timeslot p
$P_{dest}(j,k,p)$	the probability of a passenger picked up in cell
	j having a destination in cell k during timeslot p
dp(j,p)	the predicted dynamic price multiplier at cell j
	during timeslot p
f_c	the fuel consumption per kilometre
fbase	the flag-fall price
f_d	the unit price per kilometre
α	the learning rate in SARSA- λ , $0 < \alpha < 1$
γ	the discount factor of future rewards, $0 < \gamma \leq 1$
e	the probability of random exploration, $0 < \epsilon < 1$
λ	the parameter that controls the fade away speed
	of past states, $0 < \lambda < 1$

Definition 2 (Seeking Route Recommendation): Given the cell division of Beijing, GPS trajectories, RoD order data, and a subset of RoD cars Z.

Find the optimal seeking route for each car in Z to increase or maximize driver revenue. That is, for a single vacant car, when it arrives at a cell, determine which neighboring cell it should go to, or just keep seeking in the current cell. Recommendation is terminated when $\tau \ge 1440$.

Basically, we solve the seeking route recommendation problem in two steps. In the first step, we use MDP to model the environment. In MDP, the transitions between cells follow the Markov property. Between transitions, the rewards—i.e., driver earnings—are calculated based on the predicted dynamic price multipliers and the possibility of taking up passenger orders. In the second step, based on the environment, we use reinforcement learning to simulate drivers' behavior and determine the optimal seeking routes. Notations used in both steps are summarized in Table III.

A. Modeling the Environment

We use MDP to model the environment. An agent has a starting state. In every state, the agent chooses one action from many candidates and jumps to another state, getting a reward from the environment that is determined by the combination of state and action. The Markov property states that the state



Fig. 7. (a) Reverse of the incoming directions, i.e., e, and (b) actions, i.e., a.

transitions and the corresponding rewards are only dependent on the current state instead of any previous states.

In the following, we describe the states, actions, state transitions, and rewards in the MDP model.

States: We denote a state as $s = [l, p, t, e] \in S$. *l* is the index of a cell and ranges from 1 to 900. *p* is the index of current timeslot, and ranges from 1 to *N*. *t* is the number of minutes in the current timeslot. So, *p* and *t* together describe the current time.

e is defined as the reverse of the incoming direction. The incoming direction, as its name suggests, is the direction the driver arrives at the current cell during seeking. The goals of using the reverse of incoming direction are twofold. First, the whole seeking path is recorded. Second, to avoid going into a loop, we require that if a seeking driver does not pick up passengers in the last cell, he or she should not go back to the last cell after the current cell. So, the reverse of the incoming direction is ruled out as a possible choice of action. Specifically, we let $e \in E = \{ \searrow, \downarrow, \checkmark, \leftarrow, \end{cases}$ $\uparrow, \uparrow, \nearrow, \checkmark, \oslash, \emptyset$. Among these directions, \emptyset means "the driver has just dropped off a passenger and there is not a definition of incoming direction," and O means "the driver has been in this cell in the last state." For simplicity, we also denote these directions as 1-10. Fig. 7(a) illustrates these directions. In Fig. 7(a), we assume that the driver is now at the center cell, and the dashed gray arrows represent incoming directions, whereas the blue solid arrows represent the reverse of incoming directions.

As an example to explain the definition of a state, a given state $s = [250, 17, 20, \emptyset]$ means that a driver has just dropped off passengers and starts seeking at 8:20 A.M. in cell 250.

Actions: At a given state, a driver takes an action and is transitioned to the next state (e.g., jumping to a neighboring cell or keep seeking in the current cell). An action *a* could be regarded as an outgoing direction. Similarly, we let $a \in A = \{ \bigtriangledown, \uparrow, \nearrow, \rightarrow, \searrow, \downarrow, \swarrow, \leftarrow, \circlearrowright \}$, and also denote these directions as 1–9, as shown in Fig. 7(b).

Under the above definitions of e and a, it is clear that we have the following relationship between them. First, to avoid going into a loop, at a given state $s_{now} = [l_{now}, p_{now}, t_{now}, e_{now}]$, the actions $a_{unallowed}$ that satisfy the following requirements are not allowed:

$$a_{\text{unallowed}} = \begin{cases} e_{\text{now}} + 4, \text{ if } 1 \le e_{\text{now}} \le 4\\ e_{\text{now}} - 4, \text{ if } 5 \le e_{\text{now}} \le 8. \end{cases}$$
(9)

Second, if a driver takes an action a and jumps to the next state with the reverse of incoming direction e, then e = a.

State Transitions: Assuming the driver is at state $s_0 = [i, p, t, e]$, there are two different kinds of state transition after the driver takes an action *a* and jumps to cell *j*.

1) The Driver Successfully Picks Up Passengers in Cell j: This happens with a probability $P_{\text{pickup}}(j, p)$. Then, passengers are delivered to their destination, denoted by cell k, with a probability $P_{\text{dest}}(j, k, p)$. The time and distance of seeking in cell j are denoted by $t_{\text{seek}}(j, p)$ and $d_{\text{seek}}(j, p)$, respectively; and the time and distance of delivering passengers from cell j to k are denoted by $t_{\text{drive}}(j, k, p)$ and $d_{\text{drive}}(j, k, p)$, respectively. The total amount of time from cell i to cell k is $T_{\text{total},1} =$ $t_{\text{drive}}(i, j, p) + t_{\text{seek}}(j, p) + t_{\text{drive}}(j, k, p)$. After the driver delivers passengers in cell k, the state is transitioned to

$$s_1 = [k, p + (t + T_{\text{total},1})/P, (t + T_{\text{total},1})\% P, \emptyset].$$
 (10)

2) The Driver Does Not Pick Up Passengers in Cell j: This happens with a probability $1-P_{\text{pickup}}(j, p)$. The time and distance of driving from cell *i* to cell *j* are $t_{\text{drive}}(i, j, p)$ and $d_{\text{drive}}(i, j, p)$; and the time and distance of seeking in cell *j* are $t_{\text{seek}}(j, p)$ and $d_{\text{seek}}(j, p)$, respectively. The total amount of time from cell *i* to seeking in cell *j* is $T_{\text{total},2} = t_{\text{drive}}(i, j, p) + t_{\text{seek}}(j, p)$. After the driver finishes seeking in cell *j*, the state is transitioned to

$$s_2 = [j, p + (t + T_{\text{total},2})/P, (t + T_{\text{total},2})\% P, a]. (11)$$

Rewards: When a driver is transitioned between two states, a reward is obtained from the environment. We include the impact of dynamic prices into the rewards.

In the first kind of state transition, the rewards consist of two parts: 1) a positive trip fare earned by the driver and 2) a negative fuel consumption. The fuel consumption could be written as

$$R_{\text{fuel},1} = -f_c \cdot \left[d_{\text{drive}}(i,j,p) + d_{\text{seek}}(j,p) + d_{\text{drive}}(j,k,p) \right].$$
(12)

In (12), f_c is the fuel consumption per kilometer, and $d_{\text{drive}}(i, j, p) + d_{\text{seek}}(j, p) + d_{\text{drive}}(j, k, p)$ is the total distance from state S_0 to S_1 . The trip fare could be written as

$$R_{\text{trip},1} = dp(j,p) \cdot (f_{\text{base}} + f_d \cdot d_{\text{drive}}(j,k,p)).$$
(13)

In (13), f_{base} is the flag-fall price, and f_d is the unit price per kilometer. dp(j, p) is the predicted dynamic price multiplier at cell *j*, based on our dynamic price prediction model in Section IV. Specifically, dp(j, p) is measured at the location in cell *j* that has the largest number of orders in our data or the center of cell *j* if such location does not exist.

In the second kind of state transition, as the driver does not pick up passengers in cell *j*, the reward is only the negative fuel consumption in the following form:

$$R_{\text{fuel},2} = -f_c \cdot \left[d_{\text{drive}}(i,j,p) + d_{\text{seek}}(j,p) \right]. \tag{14}$$

The state transitions and rewards are illustrated in Fig. 8.



Fig. 8. State transitions and rewards.

B. Solving With SARSA- λ

Based on the environment, there are generally two different ways to design seeking routes that maximize driver revenue. The first way is dynamic programming and solves for an optimal policy. But as we use real-time dynamic price prediction, and update the parameters of the MDP model every *P* minutes, dynamic programming would become overwhelmingly complicated. Another way is reinforcement learning, and tries to obtain the optimal state–action pairs that lead to higher rewards in the long run. Comparatively, reinforcement learning is faster but still produces near-optimal results on convergence.

SARSA- λ is a typical reinforcement learning algorithm that solves for the optimal state–action pairs. SARSA stands for "State, Action, Reward, State, Action" and it works in a similar way to the famous *Q*-learning algorithm—the driver learns a *Q*-table that stores a *Q*-value for each state–action pair describing the utility of taking an action given a state, by trying different actions and observing the rewards returned from the environment. Unlike *Q*-learning, SARSA is an onpolicy approach. Based on SARSA, SARSA- λ updates the *Q*-values based on all past states and actions instead of only the current states and actions. To accomplish this, SARSA- λ introduces the eligibility trace by using an E-table and records the whole trace with which a driver has been seeking for.

Algorithm 1 shows the SARSA- λ algorithm for a single driver. The input to the algorithm is the environment modeled by MDP in Section V-A, and the dynamic price multiplier predicted by the our model in Section IV; and the output of the algorithm is a *Q*-table for the driver, so that the driver could choose the action with the highest *Q*-value at any given state. The inner loop represents the trial-and-error process of the driver through a whole day, and the outer loop is performing this process for many times until the *Q*-table converges. In the following, we explain some important lines.

- 1) *Initialization (Lines 1–3):* We first set the *Q*-table and E-table with all zeros. In addition, we also initialize the timeslot that the driver starts working as p_{start} . For example, $p_{\text{start}} = 0$ means we simulate the seeking process from 0:00, whereas $p_{\text{start}} = 17$ means the starting time is 8:00 A.M.
- 2) Initialization for the Outer Loop (Lines 5–8): In every iteration of the outer loop, we initialize the timer $\tau = 0$, the E-table with all zeros, and randomly choose a cell where the driver starts seeking.
- 3) *Exploration Versus Exploitation (Lines 9 and 14):* At any given state, the driver could either choose exploration

Algorithm 1 SARSA-λ Algorithm

Input:

1. the environment modeled by MDP, λ , ϵ , γ , and α .

2. the dynamic price prediction model in Section IV.

- **Output:** the Q-table Q(s, a) for any state-action pair.
- 1: Q(s, a) = 0 for any s and a; //Initialize Q-table.
- 2: E(s, a) = 0 for any *s* and *a*; //Initialize the eligibility trace E-table.
- 3: $p = p_{start}$; //Initialize the timeslot that the driver starts working.
- 4: while Q-table not converged do
- 5: $\tau = 0;$
- 6: E(s, a) = 0 for any s and a;
- 7: Generate a random number $1 \le l_{init} \le 900$ that represent the starting cell;
- 8: $s = (l = l_{init}, p = p_{start}, t = 0, e = 0);$ //Initialize state.
- 9: Choose action *a* at state *s* using ϵ -greedy policy from Q-table;
- 10: **while** $\tau < 1440$ **do**
- 11: Take action *a* and get the predicted dynamic price multiplier dp(j, p);
- 12: Get the new state s' and the reward $R_{s \to s'}$;
- 13: Get the new timer τ' ;
- 14: Choose action a' at state s' using ϵ -greedy policy from Q-table;
- 15: $\delta \leftarrow R_{s \rightarrow s'} + \gamma \cdot Q(s', a') Q(s, a);$
- 16: $E(s, a) \leftarrow E(s, a) + 1$; //Record the eligibility trace.
- 17: **for** all $a \in A$ and $s \in S$ **do**
- 18: $Q(s, a) \leftarrow Q(s, a) + \alpha \cdot \delta \cdot E(s, a);$
- 19: $E(s, a) \leftarrow \gamma \cdot \lambda \cdot E(s, a);$
- 20: **end for**
- 21: $s \leftarrow s', a \leftarrow a', \tau \leftarrow \tau';$
- 22: end while
- 23: end while
- 24: **return** Q-table

(i.e., randomly choosing an action) or exploitation (i.e., choosing the action with the highest *Q*-value). This is called ϵ -greedy, and ϵ is the probability of exploration. This helps avoiding being stuck at suboptimal solutions.

- 4) *Incorporating Dynamic Prices (Lines 11 and 12):* The price multiplier dp(j, p) is the result of the dynamic price prediction model. The price multiplier is then used in reward calculation.
- 5) Updating the E-Table (Lines 16–19): In the Q-learning or the original SARSA algorithm, the update of Q-value depends only on the most immediate action. In SARSA-λ, such update takes into account a series of past actions. When λ = 0, SARSA-λ reduces to SARSA; and when λ = 1, SARSA-λ remembers all the past actions. Specifically, when the driver takes an action a at a state s, the corresponding E(s, a) is added by 1, and such E(s, a) gradually fades away as time goes by.
- 6) Updating the Q-Table (Lines 15–19): The update of Q-value is on-policy. The driver takes an action a and is

transitioned from state *s* to *s'*, obtaining a reward $R_{s \rightarrow s'}$. Then, the driver again uses the ϵ -greedy policy to choose an action *a'* at state *s'*, and we calculate the TD error δ —the difference between the immediate reward plus the discounted *Q*-value of (*s*, *a*) and the current *Q*-value estimate. Both δ and E(s, a) are used to generate the new *Q*-value estimate (as in line 18).

VI. EVALUATION

We first evaluate our dynamic price prediction model, and then the reinforcement learning model.

A. Dynamic Price Prediction

1) Evaluation Metrics: We use the common metrics accuracy, precision, recall, and F1-score—to evaluate our dynamic price prediction model. These metrics are calculated based on true positive (TP), true negative (TN), false positive (FP), and false negative (FN) samples, as defined in

Accuracy =
$$\frac{TP + TN}{TP + FP + TN + FN}$$
, Preicison = $\frac{TP}{TP + FP}$
Recall = $\frac{TP}{TP + FN}$, F1 = $\frac{2}{1/Precision + 1/Recall}$. (15)

As the dynamic price prediction is a multiclass classification problem, we actually use the macro precision, macro recall, and macro F1-score in our evaluation, which are the macro average of the above metrics among all N classes.

2) Evaluation Setup: As mentioned previously, the total number of orders in our RoD order and dynamic price data is 2742120, and we generate an input feature vector for every single order. We perform two different classification tasks.

- 1) 7-*Classes Prediction:* This task has already been explained. As the price multiplier ranges from 1.0 to 1.6, it is natural to have seven classes in prediction.
- 2) 3-Classes Prediction: We also divide the price multipliers into three categories, namely, the low price multipliers (i.e., 1.0, 1.1, and 1.2), middle price multipliers (i.e., 1.3 and 1.4), and high price multipliers (i.e., 1.5 and 1.6). The goal of 3-classes prediction is to predict which category the price multiplier falls into.

We use XGBoost and LightGBM to predict dynamic price multipliers, and also implement a Random Forest and an ANN model to serve as baselines. We explain the parameters and setup of these four models below.

- ANN: There are two hidden layers that contain 256 and 128 neutrons. The drop-out regularization is added to each hidden layer to control overfitting. The activation function, output function, and loss function are ReLU, softmax, and cross-entropy, respectively. The Adam optimizer is used. The order data are divided into training set, validation set, and test set in a 7:1:2 ratio. The learning rate is set as 0.1 out of four candidates 0.01, 0.05, 0.1, and 0.5. The batch size is set as 64 out of three candidates 32, 64, and 128. These hyperparameters are selected by grid search.
- 2) *Tree-Based Models:* For tree-based models (i.e., XGBoost, LightGBM, and Random Forest), we also

Model	Hyper-parameter	Candidates and	Explanation
		the Selected Value	-
	n_estimators	[50, 100, 200 , 500]	the number of trees in the random forest
Random Forest	max_depth	[5, 10, 20, 40]	the maximum depth of a tree
Kandoni Porest	min_samples_leaf	[1, 2, 5, 10]	the minimum number of samples required to be at a leaf node
	n_estimators	[100, 500, 1000, 2000 , 4000]	the number of trees (or boosting iterations)
	learning rate	[0.01, 0.05, 0.1 , 0.5]	the learning (or shrinking) rate
LightGBM	max_depth	[2, 5, 8, 12, 20]	the maximum depth of a tree
	$\alpha_{LightGBM}$	[0.01, 0.1, 0.2, 0.5]	the parameter controlling L1-regularization
	$\lambda_{LightGBM}$	[0.01, 0.1, 0.2, 0.5]	the parameter controlling L2-regularization
	n_estimators	[100, 500, 100, 2000 , 4000]	the number of trees
	learning rate	[0.01, 0.05, 0.1 , 0.5]	the learning rate
	max_depth	[2, 5, 8, 12 , 20]	the maximum depth of a tree
XGBoost	colsample_bytree	[0.25, 0.5, 0.75, 1]	the fraction of features to be used for a tree
	subsample	[0.25, 0.5, 0.75, 1]	the fraction of instances to be used for a tree
	$\gamma_{XGBoost}$	[0.01, 0.05, 0.1 , 0.3]	the minimum split loss parameter controlling over-fitting

 TABLE IV

 MAJOR HYPERPARAMETERS IN TREE-BASED MODELS

 TABLE V

 Performances of Dynamic Price Prediction Models in Both 7-Classes and 3-Classes Tasks

Model	Accuracy		M-Precision		M-Recall		M-F1	
Model	7-classes	3-classes	7-classes	3-classes	7-classes	3-classes	7-classes	3-classes
XGBoost	0.8382	0.9067	0.8041	0.8608	0.8014	0.8464	0.8000	0.8532
LightGBM	0.7671	0.8672	0.7452	0.8121	0.6662	0.7598	0.6979	0.7812
Random	0.7941	0.8767	0.7440	0.8177	0.7579	0.7901	0.7451	0.8029
Forest								
ANN	0.6312	0.7843	0.5064	0.6805	0.3895	0.5590	0.3982	0.5858

randomly choose 20% orders as test set, and adopt a fivefold cross-validation on the remaining 80% orders. Candidates and ranges of hyperparameters are chosen based on the results and discussions of previous works (such as [40], [41], [44], and [45]), and the exact chosen values of hyperparameters are also selected by grid search from candidates and ranges, as summarized in Table IV. For each hyperparameter in Table IV, the value in bold is selected among candidates.

3) Evaluation Results: To evaluate the performance of our dynamic price prediction models, we first show the accuracy, macro precision, macro recall, and macro F1-score of XGBoost, LightGBM, Random Forest, and ANN in both the 7-classes and 3-classes prediction tasks in Table V. We have the following observations.

- XGBoost has the best performance, according to any of the four metrics and in both the 7-classes and 3-classes prediction. Specifically, in the 7-classes prediction, the XGBoost model has an accuracy of 83.82%, and a macro F1-score of 0.8000. The figures for the 3-classes prediction are 90.67% and 0.8532, respectively.
- 2) Tree-based models indeed perform better than the ANN model. This observation holds for both the 7-classes and 3-classes prediction, and even the LightGBM model, which has the weakest performance among tree-based models, is 21.53% better in accuracy than the ANN model in 7-classes prediction.
- 3) The performance improvement in 3-classes prediction compared with 7-classes prediction is much higher in the ANN model than in tree-based models. This is an interesting fact and is more obvious in macro recall

and F1-score. We hypothesize that it is the result of poor performance of ANN model on unbalanced data sets. Because dividing price multipliers into three classes means different classes are more balanced, the ANN model has a much better performance.

4) Comparing between XGBoost, LightGBM indeed has a significantly shorter training time. In our simulation with Intel i7-12700K and GeForce RTX 3070, it takes about 2500 s to converge in XGBoost, whereas the training time is only about 570 s in LightGBM.

The above results justify that XGBoost is the most suitable model to perform dynamic price prediction. It should be noted that we value prediction accuracy more than model training time, because the dynamic price prediction model could be offline and it is acceptable to update and retrain the model every half hour, every hour, or even every day.

We also evaluate the stability of Random Forest, XGBoost and LightGBM—i.e., whether a model yields similar performances when working with different data sets. To do that, we divide the RoD order and dynamic price data into fivefold, follow the similar procedures in fivefold cross validation, calculate the accuracy, macro precision, macro recall, and macro F1-score when a model works with different folds, and obtain the standard deviation of each metric across these folds. Results show that LightGBM has the largest deviation (i.e., the worst stability) for all metrics. For XGBoost and Random Forest, results are rather mixed: XGBoost has a lower standard deviation in macro precision and macro F1, whereas Random Forest has a lower deviation in accuracy and macro recall. Combined with the fact that XGBoost has the best performance, it is now safely to claim that XGBoost, among

Prediction Ground truth	1.0	1.1	1.2	1.3	1.4	1.5	1.6	Recall	F1-score
1.0	244026	3631	6522	1622	5792	634	2757	0.9209	0.9011
1.1	2914	64252	494	306	311	27	134	0.9388	0.9074
1.2	14798	2928	33702	1731	4834	231	623	0.5727	0.6437
1.3	1129	366	490	34397	1020	108	91	0.9148	0.8713
1.4	10302	1608	3979	2796	30972	1375	3479	0.5682	0.6054
1.5	242	44	47	133	491	20075	400	0.9367	0.8872
1.6	3211	353	630	368	4393	1375	32281	0.7576	0.7837
Precision	0.8822	0.8780	0.7348	0.8318	0.6478	0.8426	0.8118		

 TABLE VI

 CONFUSION MATRIX OF 7-CLASSES PREDICTION BY XGBOOST

 TABLE VII

 Confusion Matrix of 3-Classes (Low/Mid/High Price Multipliers) Prediction by XGBoost

Prediction Ground truth	low	mid	high	Recall	F1-score
low	375335	12767	4168	0.9568	0.9484
mid	19107	67493	5511	0.7327	0.7618
high	4792	4829	54422	0.8498	0.8494
Precision	0.9401	0.7932	0.8490		

all four models, is the most appropriate model for dynamic price prediction in our study.

To further evaluate the performance of the XGBoost model, Tables VI and VII show the confusion matrix of the 7-classes and 3-classes prediction by XGBoost, respectively. It is clear that as follows.

- 1) The XGBoost has a precision higher than 0.8 for five price multipliers (i.e., 1.0, 1.1, 1.3, 1.5, and 1.6), and has a recall higher than 0.9 for four price multipliers (i.e., 1.0, 1.1, 1.3, and 1.5). This indicates that the XGBoost model has a satisfactory performance for most of the price multipliers.
- 2) Results from Table VII show that the performance is the best for low price multipliers, followed by high and middle price multipliers. The corresponding F1-score are 0.9484, 0.8494, and 0.7618, respectively. A possible reason is that high and low price multipliers are more regular and predictable, and middle price multipliers are more random. It should also be noted that even with the differences of prediction accuracy among different price multipliers, our XGBoost model still has a satisfactory performance on all these price multipliers.
- 3) Results from Table VI also suggest that our model has a relatively poor performance for price multiplier 1.2 and 1.4, according to recall, precision, or F1-score. There are two possible reasons. First, 1.2 and 1.4 are price multiplies standing in the middle, and may be difficult to predict due to a high randomness, as we have already discussed. Second, it is observed from Table VI that both price multipliers are always mistakenly predicted as 1.0. This may be the result of an immature pricing algorithm, as our data set dates back to 2016 when the dynamic pricing mechanism was still at an early developing stage. So in some cases the price multipliers

are set to a wrong value, which may be contradictory to our price prediction model. To improve prediction performance, we need to either collect new data, or find out the circumstances in which the price multipliers 1.2 and 1.4 are more common and then design separate models for such circumstances.

To summarize, the XGBoost model has the best performance among the four models we implement. In the 7-classes and 3-classes prediction, the XGBoost model achieves an accuracy of 83.82% and 90.67%, respectively. Therefore, we adopt the XGBoost model to perform dynamic price prediction, and use the predicted value in our reinforcement learning model that is evaluated in Section VI-B.

B. Reinforcement Learning Model

In this section, we evaluate our reinforcing learning model that uses SARSA- λ to perform seeking route recommendation. Our main goals are twofold.

- We evaluate the effectiveness of introducing dynamic price prediction. Is it necessary to do that? What are the results if average dynamic prices are used instead or no dynamic prices are considered?
- 2) We evaluate the effectiveness of using SARSA- λ . We compare the performance of SARSA- λ with other common reinforcement learning models.

1) Evaluation Metrics: We define two metrics, namely, the revenue efficiency RE and the profit efficiency PE, to evaluate the effectiveness of seeking route recommendation. For a driver, we use R_{total} to denote the total revenue the driver makes during a specified time period; and use T_{total} to denote the total working time during this period (including the time of seeking for and delivering passengers). We also use T_{deliver} to represent the total amount of time used for delivering

passengers. Then, we have

$$RE = \frac{R_{\text{total}}}{T_{\text{total}}} \tag{16}$$

$$PE = \frac{R_{\text{total}}}{T_{\text{deliver}}}.$$
(17)

In other words, revenue efficiency RE measures the driver's revenue-making capability comprehensively, whereas profit efficiency PE focuses more on the drivers' ability to find more profitable (i.e., better) orders.

2) Evaluation Setup: We simulate our SARSA- λ approach based on our RoD order data set, RoD GPS trajectories and the dynamic price prediction model, to evaluate the effectiveness of seeking route recommendation through SARSA- λ . Across the time range of our data sets, i.e., from August to November 2016, we choose a random Friday and Saturday, as a representative weekday and weekend, to simulate our approach. The chosen days should not be a public holiday (such as the National Day holiday in China). It should also be noted that though we only present results on the chosen Friday and Saturday here, results from other weekdays and weekends show similar effects.

To achieve our main goals, we first learn the parameters of the MDP model, and then randomly choose 500 drivers who work on this Friday or Saturday. The chosen drivers should satisfy the following criteria.

- They work for at least 4 h on the chosen day, and their GPS trajectories contain few errors.
- 2) They accept more than one order on the chosen day, and these orders should be effective. For example, orders with very small trip duration or distance may be the result of inaccurate or malfunctioning GPS devices.
- 3) They also work for most of other days.

These criteria ensure that the chosen drivers have regular working patterns, and their trajectories are not outliers. We then adopt a prespecified model, simulate the trajectories of the chosen drivers, and record the orders, revenue, and trajectories of each driver. The prespecified model could be the SARSA- λ model, or other variants that serve as baselines, and they are discussed in Sections VI-B3 and VI-B4 accordingly.

Parameters of the MDP model (i.e., the environment) are set by the following procedures.

 P_{pickup}(j,p): The pickup probability in cell j during timeslot p is approximated by the ratio of the number of orders, denoted by N_{order}(j, p), to the number of vacant cars passing the cell, denoted by N_{passby}(j, p). We only count once if the same vacant car appears continually in a cell during a timeslot

$$P_{\text{pickup}}(j,p) = \frac{N_{\text{order}}(j,p)}{N_{\text{passby}}(j,p)}.$$
 (18)

P_{dest}(j, k, p): The destination probability is approximated by the historical orders. We use N_{order}(j, k, p) to record the total number of orders starting from cell j during timeslot p and ending in cell k, and then

$$P_{dest}(j,k,p) = \frac{N_{order}(j,k,p)}{N_{order}(j,p)}.$$
(19)

- 3) $t_{seek}(j, p)$ and $d_{seek}(j, p)$: The seeking distance $d_{seek}(j, p)$ in cell *j* during timeslot *p* could be either set as a fixed value—about half of the cell size—or set as a varying value in different timeslots. We try both ways and results show that the difference is small. Hence, we set $d_{seek}(j, p)$ to be 500 m. $t_{seek}(j, p)$ is then set accordingly, based on the average driving speed in cell *j* during timeslot *p*.
- 4) t_{drive}(j, k, p) and d_{drive}(j, k, p): The driving time and distance starting from cell j during timeslot p to cell k are approximated by the average driving time and distance in our historical orders. If, for some (j, k, p) combination the number of historical orders is zero, then we resort to the AMap API to check for the estimated driving time and distance instead.
- 5) f_{base} , f_d and f_c : f_{base} and f_d are set to 15 and 2.8 (both in RMB Yuan) according to the service provider's policy. f_c is set to 0.5, similar to previous studies (e.g., [27]).

Hyperparameters of the SARSA- λ model are also selected based on grid search, similar to what we do in Section VI-A2. Candidates and ranges of hyperparameters are chosen based on the discussions and suggestions in previous works, such as [35], [38], [46], and [47]. We briefly give the choice of some major hyperparameters below.

- 1) α : The learning rate α is set to 0.1 among candidates [0.03, 0.05, 0.1, 0.2, 0.3].
- 2) γ : The discount factor of future rewards γ is set to 0.5 among candidates [0.1, 0.3, 0.5, 0.7].
- 3) ϵ : The probability of random exploration ϵ is set to 0.1 among candidates [0.05, 0.1, 0.3, 0.5, 0.7].
- λ: The parameter controlling the impacts of past states *ϵ* is set to 0.5 among candidates [0, 0.1, 0.3, 0.5, 0.7, 1].

3) Effectiveness of Dynamic Price Prediction: To evaluate the effectiveness of dynamic price and dynamic price prediction, our simulation compares the SARSA- λ model with the ground-truth and two baselines.

- 1) *Real:* The ground-truth from our data.
- SL-DPP: Our SARSA-λ model explained in Section V-B. SL-DPP stands for "SARSA Lambda-Dynamic Price Prediction."
- 3) SL-ADP: SL-ADP is similar to SL-DPP. But instead of using dynamic price prediction in calculating rewards, SL-ADP uses the average historical dynamic price multiplier in a cell. SL-ADP stands for "SARSA Lambda-Average Dynamic Prices."
- SL-1.0: SL-1.0 is similar to SL-DPP, but it ignores dynamic prices at all. This is equivalent to setting all *dp(j, p)*—the price multiplier in cell *j* during timeslot *p*—to 1.0.

We first present the distribution of revenue efficiency RE, profit efficiency PE, and searching time (the amount of time used to seek for passengers) of SL-DPP, SL-ADP, and SL-1.0 on four selected time periods on weekday. The selected time periods are [8:00, 9:00), [12:00, 13:00), [18:00, 19:00), and [23:00, 0:00), and cover representative hours, such as morning and evening rush hours, night hours, and nonrush hours around noon. The distribution of revenue efficiency, profit efficiency, and searching time are shown in boxplots in

TABLE VIII Averages of Revenue Efficiency (RE), Profit Efficiency (PE), and Searching Time on Weekday

Period	Scheme	average	average	average
		RE	PE	searching time
	Real	0.88	1.11	15.49
8.00	SL-1.0	1.00	1.15	10.05
8.00	SL-ADP	1.31	1.49	9.04
	SL-DPP	1.45	1.82	14.93
	Real	0.80	1.18	22.63
12:00	SL-1.0	1.01	1.24	13.70
	SL-ADP	1.16	1.42	13.40
	SL-DPP	1.22	1.58	15.91
	Real	0.81	1.12	19.73
18.00	SL-1.0	1.01	1.22	12.09
16.00	SL-ADP	1.33	1.59	11.65
	SL-DPP	1.49	1.80	11.95
23:00	Real	0.78	1.31	28.14
	SL-1.0	1.01	1.38	18.88
	SL-ADP	1.25	1.70	18.85
	SL-DPP	1.30	1.82	20.32



Fig. 9. Distribution of revenue efficiency on selected periods on weekday.



Fig. 10. Distribution of profit efficiency on selected periods on weekday.



Fig. 11. Distribution of searching time on selected periods on weekday.

Figs. 9–11, respectively. Table VIII further gives the average values of these metrics.

Figs. 9–11, together with Table VIII, indicate that as follows.

 Using the SARSA-λ Reinforcement Learning Model Is Effective: In all four selected periods, even the



Fig. 12. Average revenue efficiency throughout the whole day (weekday).

SL-1.0 scheme improves the seeking efficiency than ground-truth. For example, during [8:00, 9:00), comparing between SL-1.0 and ground-truth, the average revenue efficiency and profit efficiency among all drivers are increased by 13.6% and 3.6%, and the average searching time is reduced by about 35%. This indicates that reinforcement learning is able to help drivers to find better orders and get orders quickly.

- 2) It Is Necessary to Consider Dynamic Prices in Seeking Route Recommendation: In all four selected periods, the revenue efficiency and profit efficiency are higher with SL-ADP and SL-DPP than with ground-truth or SL-1.0. For example, during [18:00, 19:00), the revenue efficiency in ground-truth and with SL-1.0, SL-ADP, and SL-DPP are 0.81, 1.01, 1.33, and 1.49, respectively. During [18:00, 19:00), the two schemes SL-ADP and SL-DPP that consider dynamic prices increase revenue efficiency by 31.6% and 47.5% compared to SL-1.0, respectively.
- 3) Using Predicted Dynamic Prices Further Increases Seeking Efficiency Than Simply Using Average Statistics: In all four selected periods, SL-DPP achieves higher revenue and profit efficiency than SL-ADP. Taking the revenue efficiency as an example, SL-DPP gives a revenue efficiency 10.7%, 5.2%, 12.0%, and 4.0% higher than SL-ADP does, during [8:00, 9:00), [12:00, 13:00), [18:00, 19:00), and [23:00, 0:00), respectively. We also learn from these figures that the increase of revenue efficiency in morning and evening rush hours is higher than in other two periods. Hence, dynamic price predictions prove to be effective in capturing the rapidly fluctuating dynamic prices, which is beyond the capability of using average statistics.
- 4) Reinforcement Learning Reduces Searching Time, But Considering Dynamic Prices Goes to the Opposite Direction: This is an interesting observation. It is clear from Fig. 11 that SL-DPP usually leads to a higher searching time than other schemes, whereas SL-1.0 always has a lower searching time compared with ground-truth. We consider the reason is that it takes more time for drivers to find higher price multipliers. Despite a longer searching time, drivers indeed find more profitable orders by considering dynamic prices, especially the predicted dynamic prices.

We then evaluate model performances on different hours and on weekend. Figs. 12 and 13 plot the average revenue



Fig. 13. Average revenue efficiency throughout the whole day (weekend).



Fig. 14. Average profit efficiency throughout the whole day (weekday).

efficiency throughout the day on weekday and weekend, respectively. Fig. 14 plots the average profit efficiency throughout the day on weekday. We have the following observations.

- Our above insights about the effects of using reinforcement learning, considering dynamic prices or dynamic price predictions hold throughout the day, either on weekday or weekend. In other words, SL-DPP gives the highest revenue, followed by SL-ADP, SL-1.0, and then the ground-truth.
- The amount of increase in revenue efficiency varies in different hours-of-day on weekday and weekend. For example:
 - a) SL-DPP Versus SL-ADP on Weekday: The increase is the highest during [13:00, 16:00), which is a time period with relatively stable and lower price multipliers. This shows that considering dynamic price predictions is especially important when the price multipliers across the city is lower and stable, because dynamic price prediction is an extra source of information other than the average statistics.
 - b) SL-DPP Versus SL-ADP or SL-1.0 on Weekend: The increase is the highest during [11:00, 14:00). Note that passenger demand pattern on weekend is different from that on weekday, and the number of orders peak around noon or early afternoon [38]. And according to Fig. 3, the price multiplier is relatively low and stable during this period. This observation shows again that dynamic price multiplier is helpful as an extra source of information.
- Comparing between Figs. 12 and 14 further verifies our previous observations on profit and revenue efficiency. For example:
 - a) Throughout the day, reinforcement learning helps drivers earn more by reducing the searching time,



Fig. 15. Distribution of revenue efficiency of SL-DPP, Q-DPP, and DRL-DPP on selected periods on weekday.

but the quality of orders largely remains the same. This is clear by inspecting the profit efficiency and revenue efficiency of SL-1.0 and ground-truth. From Fig. 12, it is shown that drivers indeed earn more in SL-1.0 than in ground-truth. But Fig. 14 shows that the profit efficiencies of these two schemes are very close, indicating that drivers obtain similar orders.

b) Further considering the dynamic prices, regardless of the averages or predictions, improves both driver revenue and order quality. In both Figs. 12 and 14, either SL-DPP or SL-ADP significantly increases the profit and revenue efficiency throughout the day than the ground-truth. It also proves that when considering dynamic prices, the improvement of order quality is high enough to compensate for the loss introduced by a longer searching time.

4) Effectiveness of Using SARSA- λ : Similar to Section VI-B3, to evaluate the effectiveness of using SARSA- λ , our simulation compares the SL-DPP model with the following two baselines.

- 1) *Q-DPP*: Q-DPP is similar to SL-DPP, as it uses dynamic price prediction. The difference is that Q-DPP uses *Q*-learning to recommend seeking routes instead of SARSA- λ . *Q*-learning is also a common reinforcement learning model based on *Q*-table, but it is off-policy. We also set $\epsilon = 0.1$, as in our SARSA- λ model.
- DRL-DPP: DRL-DPP is also similar to SL-DPP and Q-DPP, but it adopts a deep reinforcement learning model to recommend seeking routes, instead of SARSAλ or Q-learning. Specifically, deep Q-networks replace the Q-table: the input to the Q-network is the current state information, and the output is the Q-value. The deep reinforcement learning model estimates the Q-values by training a deep learning model, and it also adopts mechanisms such as experience replay and target Q-network to improve its performance.

We present the distribution of revenue efficiency and profit efficiency of SL-DPP, Q-DPP, and DRL-DPP on the four selected periods on weekday, in Figs. 15 and 16, respectively. We also show the average revenue efficiency and profit efficiency throughout the whole day, in Figs. 17 and 18. It is clear from these figures that as follows,

1) Compared to the Ground-Truth, SL-DPP, Q-DPP, and DRL-DPP All Significantly Increase Driver Earnings;



Fig. 16. Distribution of profit efficiency of SL-DPP, Q-DPP, and DRL-DPP on selected periods on weekday.



Fig. 17. Average revenue efficiency of SL-DPP, Q-DPP, and DRL-DPP throughout the whole day (weekday).



Fig. 18. Average profit efficiency of SL-DPP, Q-DPP, and DRL-DPP throughout the whole day (weekday).

and SL-DPP Indeed Has the Best Performance Throughout the Day: For example, during [18:00, 19:00), the average revenue efficiency of SL-DPP, Q-DPP, and DRL-DPP is 1.49, 1.39, and 1.30, respectively; and the average profit efficiency of SL-DPP, Q-DPP, and DRL-DPP is 1.80, 1.69, and 1.59, respectively. Our hypothesis is that SL-DPP performs the best because it adopts the E-table to record historical traces and this helps learning good policies. For the reason why DRL-DPP does not has a satisfactory performance, we give a more detailed discussion in Section VII.

2) The Difference Between SL-DPP and Q-DPP (or DRL-DPP) Differs Obviously During Different Hours-of-Day: For example, during the time periods with more severe supply-demand imbalance (e.g., evening rush hours), SL-DPP has a higher increase of revenue efficiency. We are inspired by this and the last observation that SL-DPP is more capable to capture the high-quality orders resulting from supply-demand imbalance and thus higher price multipliers.

Furthermore, to evaluate the effectiveness of using SARSA- λ on weekend, Fig. 19 shows the average revenue efficiency of



Fig. 19. Average revenue efficiency of SL-DPP, Q-DPP, and DRL-DPP throughout the whole day (weekend).



Fig. 20. Distribution of revenue efficiency of SL-DPP throughout the whole day (weekday versus weekend).

SL-DPP, Q-DPP, and DRL-DPP on weekend, as a comparison to Fig. 17. Fig. 20 compares the revenue efficiency of SL-DPP on weekday and weekend. Because of the limited space, we list some key observations below.

- 1) The above observations about the performance of SL-DPP, Q-DPP, and DRL-DPP also hold throughout the day on weekend.
- 2) Regarding the SL-DPP model, the revenue efficiencies on weekday and weekend show different patterns. For example, during morning or evening rush hours, the revenue efficiency is higher on weekday; during [11:00, 15:00), i.e., noon and early afternoon, the revenue efficiency is higher on weekend.

VII. SUMMARY AND DISCUSSION

We give a brief summary based on our evaluation results, and provide relevant discussions. Our study focuses on using both reinforcement learning and dynamic price prediction to provide seeking route recommendations to drivers, and we thus summarize the effects of considering dynamic price prediction, and using reinforcement learning, respectively. Furthermore, the comparison of different reinforcement learning models in Section VI-B4 gives some intriguing results, and it is necessary to include a further discussion. Finally, we discuss how dynamic prices help avoiding recommending nearby drivers to same locations, which is a common problem with single-agent reinforcement learning models.

The Effects of Considering Dynamic Price Prediction: In fact, regarding dynamic prices, we compare three different circumstances: 1) considering no dynamic prices; 2) using only the average dynamic price multipliers; and 3) using the predicted dynamic price multipliers. Among these three circumstances, our evaluation results show that the revenue efficiency increases progressively.

Besides, we also summarize the following findings from evaluation results.

- Considering dynamic prices, no matter the average statistics or the predicted price multipliers, leads to a significant improvement of order quality and thus increases driver revenue.
- 2) Considering dynamic prices, especially the predictions, increases the searching time, i.e., drivers need more time to search for better orders. But the improvement of order quality could offset the longer searching time, and thus driver revenue is further increased, compared to only using the average dynamic prices.
- 3) For busy periods with high demand, dynamic price multipliers, even the average ones, help drivers to capture the supply and demand fluctuation and thus obtain higher revenues.
- 4) For periods with generally lower price multipliers and more stable demand, the predicted dynamic price multipliers are especially important. The predicted price multipliers serve as an extra source information and help drivers to find out locations where there are better orders.

The Effects of Using Reinforcement Learning: In our study, the effects of using reinforcement learning mainly come from two aspects. First, an agent in SARSA- λ measures the utilities of taking different actions by considering long-term rewards, i.e., looking ahead. Second, SARSA- λ remembers the whole trace with the variable λ , i.e., looking backwards. By trail and error, the agent finally learns the optimal state–action pairs that maximize the expectation of the sum of rewards. Our evaluation results verify and show the following key findings.

- 1) Even without considering dynamic prices in any form, reinforcement learning alone could already help drivers to find better orders and get orders quickly. For example, our results show that during [8:00, 9:00), the average revenue efficiency is increased by 13.6%, and the average searching time is reduced by about 35%, when only SARSA- λ is used and no dynamic prices are considered.
- 2) The effects of reinforcement learning are, to some extent, in an opposite direction against the effects of dynamic prices; but combining them together yields better results. Specifically, considering dynamic prices means better orders and longer searching time, whereas using reinforcement learning means better orders and shorter searching time. Among these effects, shorter searching time is more obvious in using reinforcement learning and better orders is more obvious in considering dynamic prices. When both reinforcement learning and dynamic prices are adopted, these effects offset each other to some extent and driver revenue is further increased.

The Comparison Between Different Reinforcement Learning Models: Our evaluation results show that among the three chosen reinforcement learning models, SARSA- λ (SL-DPP) has the best performance, followed by *Q*-learning (Q-DPP), and then deep *Q*-network (DRL-DPP), and this observation holds throughout the day, either on weekday or weekend. This observation is intriguing in that deep *Q*-network has the worst performance, and in the following, we make a comparison between them.

First, deep *Q*-network is more suitable with highdimensional or continuous state space, and since our problem does not involve a high-dimensional state space, the advantage of deep *Q*-network is already diminished. Basically, for highdimensional or continuous state space, storing the *Q*-table costs a large amount of storage, making *Q*-learning or SARSA- λ unrealistic. For deep *Q*-network, only the parameters of the involved networks are stored, leading to storage space savings. In our study, we could train our SARSA- λ model for each timeslot so that the state space size is reduced, and it takes less than 2-GB memory. Therefore, the problem in our study does not require using deep *Q*-network to solve.

Second, to adopt to high-dimensional or continuous state space problems, deep *Q*-network uses neural networks to output *Q*-values. This, in fact, trades the accuracy of *Q*-values for running time and storage space.

Third, the training of deep *Q*-network is more complicated, as explained below.

- Fine-tuning requires more efforts. We find that varying some hyperparameters may lead to drastic performance degradation, and sometimes the training may not even converge. In our evaluations, we already give the best results, but it still could not outperform SARSA-λ and *Q*-learning. This also shows the difficulty of fine-tuning.
- 2) The training time is also longer. We normally train for 5000–10 000 epochs before convergence. The training time for SARSA- λ and *Q*-learning is about 10–15 min, whereas the training time for deep *Q*-network ranges between 50 min and 1 h.
- 3) The training of deep *Q*-network requires sophisticated hardware such as high-end GPUs. By comparison, only an ordinary CPU is needed in training SARSA- λ and *Q*-learning.

Finally, SARSA- λ and *Q*-learning are more light-weight, whereas deep *Q*-network is more heavy-weight. Specifically, in SARSA- λ and *Q*-learning, *Q*-values are stored as tables, and recommending seeking routes means reading from these tables, costing little memory and computation resources. To the opposite, in deep *Q*-network, either training the neural networks or calculating neural network outputs requires a lot of computation resources. Therefore, in future deployment, it would be possible to deploy the SARSA- λ or *Q*-learning models to drivers' cell phones or on-car mobile devices, but it would be highly improbable for deep *Q*-network.

The above discussions mainly focus on the advantages of SARSA- λ or *Q*-learning over deep *Q*-network, but it should be noted that deep *Q*-network has some important advantages, which may be less obvious on low or mid-dimensional problems such as ours. For example:

- 1) Deep *Q*-network could handle continuous or highdimensional state space, as mentioned previously.
- 2) By using the experience replay mechanism, deep Q-network store and reuse samples, and could thus effectively handle the data sparsity problem.
- 3) By using neural networks to approximate *Q*-values, deep *Q*-network is capable to learn complex strategies and

value functions and could thus better adapt to complex environment and tasks.

4) Deep *Q*-network supports end-to-end learning. By using deep neural networks, it is able to learn from original inputs such as images or trajectories instead of manually designed features.

To sum up, normally we may consider that deep Q-network is better than SARSA- λ or Q-learning. In fact, it would be more accurate to claim that deep Q-network could handle more complex situations and achieve a satisfactory performance in such situations. Yet when it comes to low or mid-dimensional state space problems such as ours, deep Q-network may lose its advantages in performance due to reasons such as inaccurate outputs approximated by neural networks, highly demanding fine-tuning, longer training time, etc.

The Role of Dynamic Prices in Avoiding Similar or Same Recommendations: A common problem arising in similar single-agent studies is that similar or same recommendations would be possibly made to drivers nearby (e.g., in the same cell) or drivers with similar properties. Such recommendations are undesirable, as drivers may then concentrate in a very small area, leading to a more-than-enough supply and unpredictably fluctuating dynamic price multipliers. Even multiagent reinforcement learning models are not good enough in solving this problem, as one needs a lot of parameters in order to formulate accurate models, such as drivers' emotional state, drivers' adoption rate of the recommended routes, changes to the environment (e.g., all those probabilities) after driver adoption, etc. These parameters are, unfortunately, not easy to obtain without laborious and arduous field tests and collaboration from real service providers. We thus consider multiagent models are not reliable and realistic enough until, in foreseeable future, when such field tests become possible. There are also some heuristics to tackle problem, such as generating a list of possible recommendations (that are equally or almost equally optimal) and randomly picking one after another for drivers nearby.

The adoption of dynamic prices in RoD service offers a new perspective in solving this problem. For example, if the pricing algorithm is perfectly designed and could respond to the changes of supply and demand in real time, then the price multiplier within a small "hot" area would go down when drivers begin to gather there. This lowers the rewards of going to this hot area, and keeps other drivers away. With the adoption of dynamic prices, even single-agent models such as ours could avoid recommending a particular cell to a lot of drivers. In our study, on the one hand, we assume that the dynamic pricing mechanism of the service provider from which we obtain our data sets satisfies the above requirement, but how to design such a mechanism is another story and is out of the scope of this article. On the other hand, our inclusion of a dynamic price prediction model could be regarded as a way to imitate the service provider's dynamic pricing mechanism: such model provides a description of the environment at a finer spatial and temporal granularity, compared with the traditional descriptions by the pickup probabilities or destination probabilities.

VIII. CONCLUSION

We focus on the seeking route recommendation problem, i.e., recommending the next cell to a seeking driver so that driver revenue is higher, in RoD service. Though this problem has been studied from various perspectives in taxi service, RoD has two new features—dynamic pricing and data-driven—that enable us to improve driver revenue to the next level.

By analyzing real service data, we point out that it is necessary to both recommend seeking routes to drivers and take into account dynamic price multipliers in generating recommendations. We first design a dynamic price prediction model to generate the predicted price multiplier given the time and location described by features from multisource urban data. We then adopt a reinforcement learning model, and calculate the rewards of transitions between cells based on the predicted price multipliers.

Evaluation results first validate the effectiveness of both models. The dynamic price prediction model achieves a satisfactory accuracy of 83.82% and 90.67%, in the 7-classes prediction (i.e., predicting the exact multiplier) and 3-classes prediction (i.e., predicting whether the multiplier is high, mid, or low), respectively. The reinforcement learning model also significantly increases drivers' average revenue as well as profit efficiency and reduces average searching time. We also emphasize the positive effects of considering dynamic price predictions: using predicted price multipliers is better than using average price multipliers, which is better than considering no dynamic prices at all, in terms of both revenue and profit efficiency.

Besides effectiveness validation, our results also reveal some interesting facts. For example, using reinforcement learning primarily reduces searching time, whereas considering dynamic prices as well as predictions mainly enables drivers to improve order quality, though meanwhile the searching time becomes longer. This shows the different roles reinforcement learning and dynamic prices play in increasing driver revenue. We also find out that dynamic price prediction has positive effects because it serves as a new and reliable source of information and captures the supply and demand fluctuation. Furthermore, by comparing between multiple reinforcement learning models, simple models, such as SARSA- λ and *Q*-learning, prove to have better performance than complex models such as deep reinforcement learning.

For future work, we primarily consider the application and implementation on real service. Examples include autonomous driving and implementing our models on on-car low power platforms that could automatically give suggestions to drivers throughout working shifts. But there are many details and perspectives to consider before doing all these. One of the most important considerations in practical applications may be the use of multiagent models to comprehensively depict the picture of a large number of drivers, and it is important to first obtain a lot more information such as the adoption ratio of recommended routes, the changes of passenger patterns after observing new driver behavior, the behavior patterns of drivers, etc. As an example, Ghosh et al. [48], [49] adopted an algorithm based on multiarmed bandit to discuss the competition and cooperation between multiple agents. These works inspire us to leave the interaction between drivers as our future work. We are currently working actively to push our collaboration with service providers to obtain the necessary information and data sets.

ACKNOWLEDGMENT

The authors thank the anonymous editors and reviewers for their constructive feedback and comments.

REFERENCES

- [1] A. Picchi. "Uber vs. taxi: Which is cheaper?" 2016. [Online]. Available: http://bit.ly/2DMgrMc
- [2] Y. M. Nie, "How can the taxi industry survive the tide of ridesourcing? evidence from Shenzhen, China," Transp. Res. Part C, Emerg. Technol., vol. 79, pp. 242-256, Jun. 2017.
- [3] L. Rayle, D. Dai, N. Chan, R. Cervero, and S. Shaheen, "Just a better taxi? a survey-based comparison of taxis, transit, and ridesourcing services in San Francisco," Transp. Policy, vol. 45, pp. 168-178, Jan. 2016.
- [4] J. D. Hall, C. Palsson, and J. Price, "Is Uber a substitute or complement for public transit?" J. Urban Econ., vol. 108, pp. 36-50, Nov. 2018. [Online]. Available: https://bit.ly/2K6Vs7L
- T. Berger, C. Chen, and C. B. Frey, "Drivers of disruption? Estimating the Uber effect," Eur. Econ. Rev., vol. 110, pp. 197-210, Nov. 2018.
- [6] Y. Cao, T. S. Gruca, and B. R. Klemz, "Internet pricing, price satisfaction, and customer satisfaction," Int. J. Electron. Commer., vol. 8, no. 2, pp. 31-50, 2003.
- [7] M. L. Kasavana and A. J. Singh, "Online auctions: Dynamic pricing and the lodging industry," J. Hosp. Leisure Market., vol. 9, nos. 3-4, pp. 127–140, 2001.
- [8] M. D. Wittman and P. P. Belobaba, "Dynamic pricing mechanisms for the airline industry: A definitional framework," J. Revenue Pricing Manag., vol. 18, pp. 100-106, Apr. 2019.
- [9] J. M. Betancourt, A. Hortaçsu, A. Oery, and K. R. Williams, "Dynamic price competition: Theory and evidence from airline markets," Nat. Bureau Econ. Res., Cambridge, MA, USA, Working Paper 30347, Aug. 2022. [Online]. Available: http://www.nber.org/papers/w30347
- [10] O. Besbes and A. Zeevi, "Dynamic pricing without knowing the demand function: Risk bounds and near-optimal algorithms," Oper. Res., vol. 57, no. 6, pp. 1407-1420, 2009.
- [11] J. Hall, C. Kendrick, and C. Nosko. "The effects of Uber's surge pricing: A case study." Oct. 2015. [Online]. Available: http://bit.ly/2kayk90
- [12] J. Gan, B. An, H. Wang, X. Sun, and Z. Shi, "Optimal pricing for improving efficiency of taxi systems," in Proc. 22nd Int. Joint Conf. Artif. Intell., 2013, pp. 2811-2818.
- [13] L. Rayle, S. Shaheen, N. Chan, D. Dai, and R. Cervero. "Appbased, on-demand ride services: Comparing taxi and ridesourcing trips and user characteristics in San Francisco," 2014. [Online]. Available: http://bit.ly/2kVkahg
- [14] L. Chen, A. Mislove, and C. Wilson, "Peeking beneath the hood of Uber," in Proc. ACM Conf. Internet Meas. Conf., 2015, pp. 495-508.
- [15] S. Guo, C. Chen, Y. Liu, K. Xu, and D. M. Chiu, "Modelling passengers' reaction to dynamic prices in ride-on-demand services: A search for the best fare," Proc. ACM Interact. Mobile Wearable Ubiquitous Technol., vol. 1, no. 4, pp. 1-23, 2018.
- [16] S. Guo, Y. Liu, K. Xu, and D. M. Chiu, "Understanding ride-ondemand service: Demand and dynamic pricing," in Proc. IEEE Int. Conf. Pervasive Comput. Commun. Workshops (PerCom Workshops), 2017, pp. 509-514.
- [17] H. Chen et al., "InBEDE: Integrating contextual bandit with TD learning for joint pricing and dispatch of ride-hailing platforms," in Proc. IEEE Int. Conf. Data Min., 2019, pp. 61-70.
- [18] M. K. Chen, "Dynamic pricing in a labor market: Surge pricing and flexible work on the Uber platform," in Proc. ACM Conf. Econ. Comput., 2016, pp. 455–455.
- [19] P. Cohen, R. Hahn, J. Hall, S. Levitt, and R. Metcalfe. "Using big data to estimate consumer surplus: The case of Uber." 2016. [Online]. Available: http://bit.ly/2pqXiWo

- [20] S. Guo, C. Chen, Y. Liu, K. Xu, and D. M. Chiu, "It can be cheaper: Using price prediction to obtain better prices from dynamic pricing in ride-on-demand services," in Proc. 14th EAI Int. Conf. Mobile Ubiquitous Syst. Comput., Netw. Serv., 2017, pp. 146-155.
- [21] S. Guo, C. Chen, J. Wang, Y. Liu, K. Xu, and D. M. Chiu, "Finegrained dynamic price prediction in ride-on-demand services: Models and evaluations," Mobile Netw. Appl., vol. 25, pp. 505-520, Apr. 2020.
- [22] S. Guo, C. Chen, J. Wang, Y. Liu, K. Xu, and D. M. Chiu, "Dynamic price prediction in ride-on-demand service with multi-source urban data," in Proc. 15th EAI Int. Conf. Mobile Ubiquitous Syst. Comput., Netw. Services, 2018, pp. 412-421.
- [23] S. Guo, C. Chen, Y. Liu, K. Xu, B. Guo, and D. M. Chiu, "How to pay less: A location-specific approach to predict dynamic prices in ride-ondemand services," IET Intell. Transp. Syst., vol. 12, no. 7, pp. 610-618, 2018.
- [24] S. Guo et al., "A simple but quantifiable approach to dynamic price prediction in ride-on-demand services leveraging multi-source urban data," Proc. ACM Interact. Mobile Wearable Ubiquitous Technol., vol. 2, no. 3, pp. 1-24, 2018.
- [25] B. Li et al., "Hunting or waiting? Discovering passenger-finding strategies from a large-scale real-world taxi dataset," in Proc. IEEE Int. Conf. Pervasive Comput. Commun. Workshops (PerCom Workshops), 2011, pp. 63-68.
- [26] D. Zhang et al., "Understanding taxi service strategies from taxi GPS traces," IEEE Trans. Intell. Transp. Syst., vol. 16, no. 1, pp. 123-135, Feb. 2015.
- [27] S. Guo et al., "ROD-Revenue: Seeking strategies analysis and revenue prediction in ride-on-demand service using multi-source urban data," IEEE Trans. Mobile Comput., vol. 19, no. 9, pp. 2202-2220, Sep. 2020.
- [28] 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 Proc. 25th ACM Int. Conf. Inf. Knowl. Manag., 2016, pp. 2329-2334.
- [29] X. Yu, S. Gao, X. Hu, and H. Park, "A Markov decision process approach to vacant taxi routing with e-hailing," Transp. Res. Part B, Methodol., vol. 121, pp. 114-134, Mar. 2019.
- [30] X. Zhou et al., "Optimizing taxi driver profit efficiency: A spatial network-based Markov decision process approach," IEEE Trans. Big Data, vol. 6, no. 1, pp. 145-158, Mar. 2020.
- [31] Z. Shou, X. Di, J. Ye, H. Zhu, and R. Hampshire, "Where to find next passengers on e-hailing platforms?-A Markov decision process approach," 2020, arXiv:1905.09906.
- [32] C.-M. Tseng, S. C.-K. Chau, and X. Liu, "Improving viability of electric taxis by taxi service strategy optimization: A big data study of New York city," IEEE Trans. Intell. Transp. Syst., vol. 20, no. 3, pp. 817-829, Mar. 2019.
- [33] N. Garg and S. Ranu, "Route recommendations for idle taxi drivers: Find me the shortest route to a customer!" in Proc. 24th ACM SIGKDD Int. Conf. Knowl. Discov. Data Min., 2018, pp. 1425-1434.
- [34] Y. Gao, D. Jiang, and Y. Xu, "Optimize taxi driving strategies based on reinforcement learning," Int. J. Geogr. Inf. Sci., vol. 32, no. 8, pp. 1677-1696, 2018.
- [35] M. Han, P. Senellart, S. Bressan, and H. Wu, "Routing an autonomous taxi with reinforcement learning," in Proc. 25th ACM Int. Conf. Inf. Knowl. Manag., 2016, pp. 2421-2424.
- [36] C. Yan, H. Zhu, N. Korolko, and D. Woodard, "Dynamic pricing and matching in ride-hailing platforms," Naval Res. Logist., vol. 67, no. 8, pp. 705-724, 2020.
- [37] H. A. Chaudhari, J. W. Byers, and E. Terzi, "Putting data in the driver's seat: Optimizing earnings for on-demand ride-hailing," in Proc. 11th ACM Int. Conf. Web Search Data Min., 2018, pp. 90-98.
- [38] S. Guo et al., "Seeking based on dynamic prices: Higher earnings and better strategies in ride-on-demand services," IEEE Trans. Intell. Transp. Syst., vol. 24, no. 5, pp. 5527-5542, May 2023.
- [39] "API of AMap service." AMap. 2023. [Online]. Available: https://lbs.amap.com/
- [40] L. Grinsztajn, E. Oyallon, and G. Varoquaux, "Why do tree-based models still outperform deep learning on typical tabular data?" in Proc. 36th Conf. Neural Inf. Process. Syst., 2022, pp. 507–520. [41] R. Shwartz-Ziv and A. Armon, "Tabular data: Deep learning is not all
- you need," Inf. Fusion, vol. 81, pp. 84-90, May 2022.
- [42] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discov. Data Min., 2016, pp. 785-794.
- [43] G. Ke et al., "LightGBM: A highly efficient gradient boosting decision tree," in Proc. 31st Conf. Neural Inf. Process. Syst., 2017, pp. 1-9.

- [44] C. Bentéjac, A. Csörgő, and G. Martínez-Muñoz, "A comparative analysis of gradient boosting algorithms," Artif. Intell. Rev., vol. 54, no. 3, pp. 1937-1967, 2021.
- [45] X. Shi, Y. D. Wong, M. Z.-F. Li, C. Palanisamy, and C. Chai, "A feature learning approach based on XGBoost for driving assessment and risk prediction," Accid. Anal. Prev., vol. 129, pp. 170-179, Aug. 2019.
- [46] H. Anas, W. H. Ong, and O. A. Malik, "Comparison of deep Q-Learning, Q-Learning and SARSA reinforced learning for robot local navigation," in Proc. 9th Int. Conf. Robot Intell. Technol. Appl. 6, 2022, pp. 443-454.
- [47] I. Zamora, N. G. Lopez, V. M. Vilches, and A. H. Cordero, "Extending the OpenAI Gym for robotics: A toolkit for reinforcement learning using ROS and Gazebo," 2017, arXiv:1608.05742.
- [48] A. Ghosh, A. Sankararaman, K. Ramchandran, T. Javidi, and A. Mazumdar, "Competing bandits in non-stationary matching markets," IEEE Trans. Inf. Theory, vol. 70, no. 4, pp. 2831-2850, Apr. 2024.
- [49] A. Ghosh, A. Sankararaman, K. Ramchandran, T. Javidi, and A. Mazumdar, "Decentralized competing bandits in non-stationary matching markets," 2022, arXiv:2206.00120.



Jintao Ke received the Ph.D. degree from Hong Kong University of Science and Technology, Hong Kong, in 2020.

He is an Assistant Professor with the Department of Civil Engineering, University of Hong Kong, Hong Kong. His research interests include shared mobility on demand, transportation big data analytics, multimodal intelligent transportation systems, transportation pricing, and short-term travel demand forecasting.



Jingyuan Wang received the Ph.D. degree from Tsinghua University, Beijing, China, in 2011.

He is currently a Full Professor with Beihang University, Beijing. His general area of research is data mining and machine learning, with special interests in smart cities.



Suiming Guo received the Ph.D. degree from the Chinese University of Hong Kong, Hong Kong, in 2018.

He is currently an Associate Professor with the College of Information Science and Technology, Jinan University, Guangzhou, China. His research interests include data mining, urban computing, pervasive computing, and smart cities studies.



Saiqin Long received the Ph.D. degree from the South China University of Technology, Guangzhou, China, in 2014.

She is currently a Professor with the College of Information Science and Technology, Jinan University, Guangzhou, China. Her research interests include cloud computing, edge computing, parallel and distributed systems, and Internet of Things.



Baoying Deng is currently pursuing the master's degree with the College of Information Science and Technology, Jinan University, Guangzhou, China. His research interests include urban computing and smart cities studies.





Chao Chen received the Ph.D. degree from UMPC (Paris 6), Paris, France, and Telecom SudParis, Évry, France, in 2014.

He is currently a Full Professor of Computer Science with Chongqing University, Chongqing, China. His research interests include pervasive computing, social network analysis, and mobile crowdsensing.



Ke Xu (Fellow, IEEE) received the Ph.D. degree from Tsinghua University, Beijing, China, in 2001. He is currently a Full Professor with the Department of Computer Science and Technology, Tsinghua University. His research interests include next-generation Internet, P2P systems, Internet of Things, network virtualization, and optimization.