Sensing the Pulse of Urban Refueling Behavior: A Perspective from Taxi Mobility

FUZHENG ZHANG, University of Science and Technology of China
NICHOLAS JING YUAN, Microsoft Research
DAVID WILKIE, University of North Carolina at Chapel Hill
YU ZHENG and XING XIE, Microsoft Research

Urban transportation is an important factor in energy consumption and pollution, and is of increasing concern due to its complexity and economic significance. Its importance will only increase as urbanization continues around the world. In this article, we explore drivers' refueling behavior in urban areas. Compared to questionnaire-based methods of the past, we propose a complete data-driven system that pushes towards real-time sensing of individual refueling behavior and citywide petrol consumption. Our system provides the following: detection of individual refueling events (REs) from which refueling preference can be analyzed; estimates of gas station wait times from which recommendations can be made; an indication of overall fuel demand from which macroscale economic decisions can be made, and a spatial, temporal, and economic view of urban refueling characteristics. For individual behavior, we use reported trajectories from a fleet of GPS-equipped taxicabs to detect gas station visits. For time spent estimates, to solve the sparsity issue along time and stations, we propose context-aware tensor factorization (CATF), a factorization model that considers a variety of contextual factors (e.g., price, brand, and weather condition) that affect consumers' refueling decision. For fuel demand estimates, we apply a queue model to calculate the overall visits based on the time spent inside the station. We evaluated our system on large-scale and real-world datasets, which contain 4-month trajectories of 32,476 taxicabs, 689 gas stations, and the self-reported refueling details of 8,326 online users. The results show that our system can determine REs with an accuracy of more than 90%, estimate time spent with less than 2 minutes of error, and measure overall visits in the same order of magnitude with the records in the field study.

Categories and Subject Descriptors: H.2.8 [**Database Management**]: Database Applications—*Spatial databases and GIS*

General Terms: Design, Implementation Algorithms, Experimentation, Performance

Additional Key Words and Phrases: Refueling event, spatial-temporal unit, expected duration, arrival rate

ACM Reference Format:
Fuzheng Zhang, Nicholas Jing Yuan, David Wilkie, Yu Zheng, and Xing Xie. 2015. Sensing the pulse of urban refueling behavior: A perspective from taxi mobility. ACM Trans. Intell. Syst. Technol. 6, 3, Article 37 (April 2015), 23 pages.
DOI: http://dx.doi.org/10.1145/2644828

1. INTRODUCTION

A growing share of the global population currently lives in cities, indicating that the trend of urbanization has become increasingly significant. With this in mind, urban

---

Authors' addresses: F. Zhang, University of Science and Technology of China, He Fei, China; email: zhfzh@mail.ustc.edu.cn; N. J. Yuan, Y. Zheng, and X. Xie, Microsoft Research, Beijing, China; emails: {nicholas.yuan, yuzheng, xingx}@microsoft.com; D. Wilkie, University of North Carolina at Chapel Hill, Chapel Hill, North Carolina; email: wilkie@cs.unc.edu.

DOI: http://dx.doi.org/10.1145/2644828

transportation issues are of foremost importance to support the individual and freight mobility requirements of large urban agglomerations. However, transportation authorities rarely have access to a real-time view of traffic statuses or patterns. Additionally, due to the heavy and growing reliance on petroleum and the environmental impact of emissions from fossil fuels, energy consumption from urban transportation represents a pressing challenge. An integral and underresearched component of the transportation system is the refueling behavior of individual cars.

An understanding of consumers' refueling activities provides the knowledge to determine the effect of different network configurations of retail fuel outlets, which is especially helpful for fuel marketing strategy and government planning. However, to a great extent, refueling behavior itself has been constantly neglected in previous empirical and theoretical studies of travel behavior and fuel use. Even when refueling is the sole purpose of a trip, it can be classified as either a shopping or personal business trip. More generally, when it is part of a multipurpose trip, the refueling stop is presumably ignored and the trip legs before and after refueling are treated as one continuous trip [Memmott 1963]. Several traditional survey-based methods related to refueling preference [Kitamura and Sperling 1987; Kelley and Kuby 2013] are limited to the self-reports from small-scale subjects, which are insufficient to represent the condition of the whole city. In addition, energy use in vehicle transportation is difficult to ascertain. This is especially true for real-time estimates. On the one hand, gas stations are typically owned by an assortment of different, competing organizations that do not want to make data available to competitors. There is also a cost associated with monitoring and publicizing data, from which station owners would derive no benefit. On the other hand, estimating energy use is also a difficult problem, as it is a function of the cars' acceleration, which is highly variable and difficult to estimate.

Given the preceding information, we propose a complete data-driven framework to understand urban refueling activity. We focus on identifying refueling events (REs), estimating the time spent and number of consumers in each spatial-temporal unit (the division of spatial-temporal space will be detailed in Section 2), and compare the refueling patterns among different social groups (including taxi drivers, household customers, and urban results as another group). Our approach uses a "human as a sensor" approach by analyzing and drawing inferences from GPS trajectories passively collected by taxicabs. First, we take a novel approach to the detection of REs, which are visits by taxis to gas stations. The detection of REs includes the time spent waiting at the gas station and the time spent refueling the vehicle. For units that cover enough detected REs, the related time spent is estimated from the distribution of the REs. For those units with few or even without REs, our context-aware tensor factorization (CATF) model will predict the time spent, which uses a context-aware collaborative filtering approach by considering various contextual features that affect consumers' refueling decisions. Further on, we treat each gas station as a queue system, and time spent in the station is used to calculate the number of customers in that unit. Therefore, the output is a global estimate of time spent and fuel use at each gas station during each time period. Finally, we compare the refueling patterns among different social groups in terms of their spatial, temporal, and economic aspects.

Our evaluation consists of the following parts. First, we conducted several experiments on the performance of our RE detection algorithm. We analyzed the recall and precision of our method on a manually labeled GPS dataset as well as a dataset generated by the authors. Next, we presented the performance of time-spent estimation and the effectiveness of the CATF model. Finally, we evaluated the effectiveness of arrival rate by comparing the number of customers calculated in the queue model with the results collected in a case study.

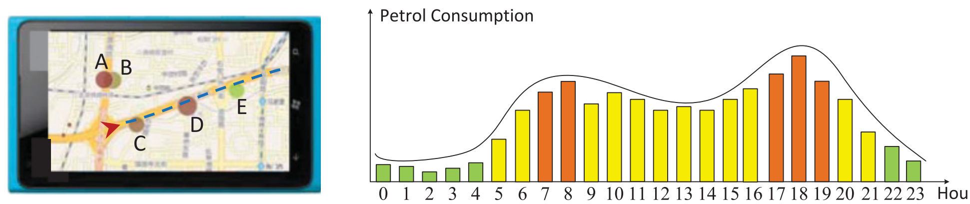

(a) Refueling recommendation (b) Petrol consumption

Fig. 1. Application scenarios of refueling activity understanding.

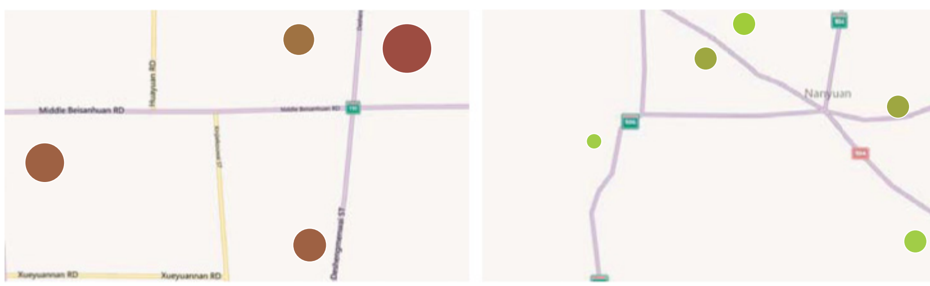

(a) Lack of demand (b) Demand surplus

Fig. 2. A local view of gas stations.

The results of our system can be applied primarily in the following scenarios:

—*Refueling recommendation*: Figure 1(a) shows several gas stations' time spent at a point in time. The darker the color, the more time spent. Assuming that a driver is in the position denoted by the arrow, even if station C is the closest, other stations might be recommended due to a shorter wait times (i.e., station E is a satisfactory choice).
—*Gas station planning*: Figure 2 shows a local view of several gas stations. The size of the stations indicates the average visit rate. The larger the size, the more drivers have visited that station. We see the number of vehicles that have refueled in the area shown in Figure 2(a), and thus these stations have a long wait time (colored red). It might be worthwhile to consider building a new gas station nearby to relieve the issue of insufficient supply. On the contrary, Figure 2(b) indicates that the gas stations are very dense in this area even though very few drivers have visited (size is small). Therefore, the government could consider closing some of them to reduce waste.
—*Energy consumption analysis*: In Figure 1(b), the curve gives a direct view of this city's time-varying petrol consumption, based on the drivers' arrival rate during each period. This can be used by station operators to formulate better commercial strategies.

This article offers a step toward persistent monitoring of urban transportation energy use and refueling behavior by passive human participatory sensing. The main contributions include the following:

—We propose a method for the discovery of REs from raw GPS trajectories.
—We analyze diverse factors (spatial, temporal, economic, climate, etc.) that influence consumers' refueling choice and incorporate them into a context-aware collaborative filtering method to estimate the time spent.
—We develop an approach that uses a queue system to calculate the overall arrival rate at gas stations from the time spent during a period.
—We uncover and compare different social groups' refueling preference in spatial, temporal, and economic aspects.

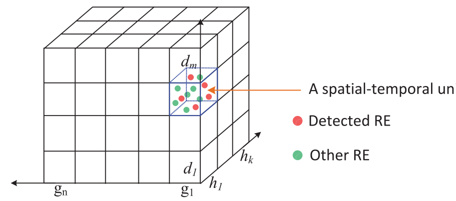

Fig. 3. The spatial-temporal division of REs.

—We evaluate our system using large-scale and real-world datasets, which consist of a trajectory dataset, a point-of-interest (POI) dataset, a road network dataset, and refueling consumption details from the reports of tens of thousands of online users.

The rest of this article is structured as follows. In Section 2, we clarify several terms widely used in this work. In Section 3, we give a detailed description of the datasets. In Section 4, we give an overview of the main components in our intelligent system. In Section 5, we discuss how to discover REs from raw trajectories. In Section 6, we discuss how to estimate time spent for each unit, especially with the aid of various contextual features to solve the sparsity issue. In Section 7, we discuss how to estimate the number of overall visits based on the time spent in the gas station. In Section 8, we evaluate the performance of our system on various datasets. In Section 9, we give a spatial, temporal, and economic view of the refueling behavior of different groups. In Section 10, we discuss the generalization and limitations of this work. In Section 11, related work is discussed and compared with our system. Finally, the conclusion and future research possibilities are presented in Section 12.

2. PRELIMINARY

In this section, we will clarify some terms used in this work.

Trajectory. A trajectory is a sequence of GPS points that is composed of a latitude, a longitude, and a timestamp.

Point of interest. A POI refers to a specific point location that someone may find useful or interesting. It is described by a latitude, a longitude, and a category (e.g., restaurant or gas station).

Refueling event. An RE refers to a vehicle refueling at a gas station. It is composed of the arrival time, departure time, and the selected gas station. An RE's duration represents the time spent there, which is the difference between the arrival time and departure time.

Spatial-temporal unit. A spatial-temporal unit is a division for REs. A unit U_{ijk} corresponds to a gas station g_i with the timestamp of hour h_j and the timestamp of day d_k, as shown in Figure 3. Each RE falls in one certain unit (its selected gas station is mapped to g_i, and its arrival time is mapped to h_j and d_k). A unit is the finest granularity that we use for refueling behavior analysis. For each unit, there are two indicators with which we are concerned: expected duration and arrival rate. The expected duration refers to how much time, on average, is spent by the vehicles refueling in this unit. The arrival rate indicates how many drivers fall in this unit.

3. DATASET

Our system is built on four kinds of data sources.

Taxi trajectories. This dataset covers GPS trajectories collected by 32,476 taxicabs located in Beijing during the period from August to November 2012, for a total of 112 days. The sampling interval of the GPS devices is about 1.01 minutes, and on average, a taxi travels 231.29km per day.

(a) Refuel information (b) Station description

Fig. 4. Self-reports collection procedure.

Road network. The road network dataset covers the road network of Beijing, which contains 106,579 road nodes and 141,380 road segments in the urban area, and in which each segment is described as a sequence of geospatial points as well as some other attributes (road level, the number of lanes, etc.).

Points of interest. There are 369,668 POIs with 602 categories. Gas stations are one category of particular interest. There are 1,221 gas stations within this city, of which 689 gas stations are located in the areas covered by our road network. In our system, we only concentrate on the 689 stations.

Self-reports. LiCheWang[1] is a vehicle-usage experience Web site where users can share their vehicle-related fees, such as fuel consumption, maintenance costs, and tickets. We collected the self-reported refueling details of 8,326 users in Beijing from this Web site. We first extracted users who live in Beijing according to their profile description. Then we crawled these users' refueling information as shown in Figure 4(a). For each refueling record, the selected gas station's description is shown in Figure 4(b). We match this station to the station in our POI dataset by the location information represented as latitude and longitude. Each record represents a refueling experience of one user, which includes the selected gas station, the day timestamp, and the refueling expenditure.

4. SYSTEM OVERVIEW

Our system provides insight into urban refueling behavior. This behavior is captured by estimating each spatial-temporal unit's expected duration and arrival rate and comparing different social groups' refueling patterns. As a preliminary step, we first identify the REs in the trajectory data. Then, for units with a sufficient number of REs, we model the expected duration as the average of the values of the contained REs. For units that have few or even no REs, we consider various factors influencing drivers' refueling decision and propose CATF to predict the expected durations. With the expected durations estimated, we utilize the queuing theoretic model to calculate each unit's arrival rate. Finally, with each unit's expected duration and arrival rate estimated, we can compare different social groups' refueling patterns from different perspectives. There are four main components in our system, as shown in Figure 5.

Refueling event detection. In this component, a large number of candidates are extracted from raw trajectories and then a filtering algorithm is applied to obtain the final results.

Expected duration learning. For a unit containing sufficient detected REs, its expected duration is represented with the detected REs' average durations. Then, for units with insufficient REs, the CATF method considers various factors that affect the

[1] http://www.liche365.com.

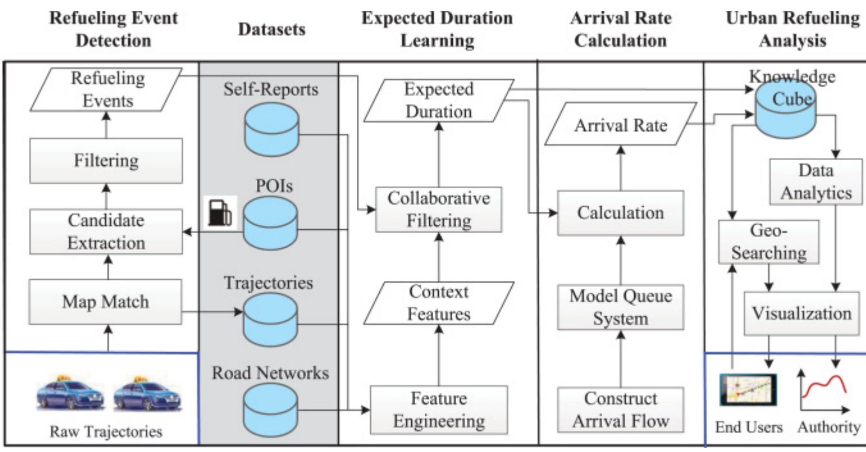

Fig. 5. System overview.

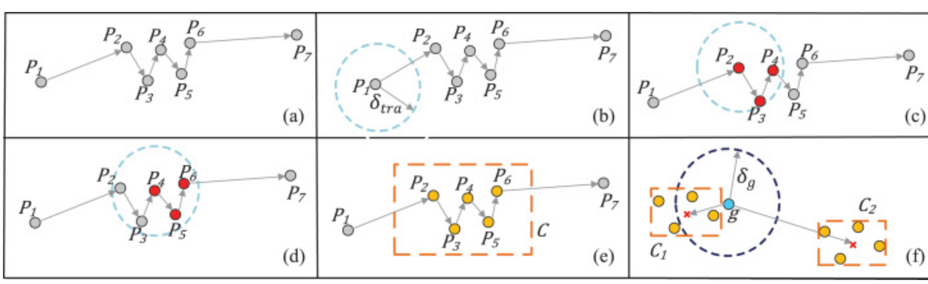

Fig. 6. Candidate extraction.

refueling decision and train a context-aware collaborative filtering method to predict their expected durations.

Arrival rate calculation. On top of the queuing theoretic model, we make a statistical inference of its arrival rate depending on a unit's expected duration.

Urban refueling analysis. According to the previous data acquired, we analyze different groups' refueling patterns in terms of spatial, temporal, and economic aspects.

5. REFUELING EVENT DETECTION

By mapping the geospatial movement of cars to the positions of gas stations, it seems that REs can easily be discovered. However, this direct approach has difficulties due to the noise of the GPS readings, and it cannot support perfect matching. The GPS devices generally have an error of 10m, and the position of a gas station is merely depicted as a single point (which is actually an area with hundreds of square meters): these two factors lead the direct approach to mistake other behavior for REs while passing up real refueling behavior. This section details the process of detecting REs from the taxis' raw trajectories when there is a certain amount of uncertainty. We first extract the refueling candidates and then use a supervised method to filter the errant candidates.

5.1. Candidate Extraction

We extract RE candidates with a consideration of mobility and geographic constraints.

For the mobility constraint, an RE corresponds to a period of slow movement. As shown in Figure 6(a), given a trajectory $p_1 \rightarrow p_2 \rightarrow \cdots \rightarrow p_7$, we first keep on checking

(a) Real RE (b) Traffic jam or waiting signal (c) Parking place

Fig. 7. Real RE with respect to pseudocandidates.

the distance between the current point and later point until it is larger than a threshold δ_{tra}. As shown in Figure 6(b), since $dist(p_1, p_2) > \delta_{tra}$, we move next and take p_2 as pivot point. We find that $dist(p_2, p_3) < \delta_{tra}$ and $dist(p_2, p_4) < \delta_{tra}$ while $dist(p_2, p_5) > \delta_{tra}$, as shown in Figure 6(c). If the interval between $p_2.t$ and $p_4.t$ is smaller than τ, p_2, p_3, p_4 forms a small cluster. Then, as shown in Figure 6(d), we fix p_4 as a pivot point to check on the later points. Finally, we take p_2, p_3, p_4, p_5, p_6 as a stay point, which is shown in Figure 6(e).

For the geographic constraint, we check the distance between a stay point's center point and the nearest gas station. Then we take those stay points that satisfy $dist(c, g) < \delta_g$ as possible RE candidates. As shown in Figure 6(f), C_1 is kept and C_2 is discarded.

We manually labeled 200 real REs by plotting the raw trajectories in digital maps and used this dataset to set the parameters (δ_{tra}, τ, and δ_g) with a grid search.

5.2. Filtering

The candidate extraction process finds clusters of points in close proximity to gas stations. However, a candidate could be generated by some other behavior. For example, for gas stations that are close to roads or intersections, the candidate might imply a traffic jam or a car waiting for signals at a traffic light. Some other POIs, such as repair shops, car washes, or even parking lots, might be located close to gas stations and create spurious candidates. Figure 7 shows a real RE compared to pseudocandidates. To filter these non-REs out of the candidate pool, we apply a supervised model using spatial-temporal and POI features.

Spatial-temporal features include the following: (1) *encompassment*, a binary value indicating whether the gas station is contained in the candidate's minimum bounding box; (2) *gas station distance*, the average distance between the candidate's points and the gas station; (3) *distance to road*, the average distance between the candidate's points and a matched road segment; (4) *minimum bounding box ratio*, the ratio between the minimum bounding box's width and height represented as $Min(\frac{width}{length}, \frac{length}{width})$; and (5) *duration*, the temporal duration of a candidate.

POI features include the following: (1) *neighbor count*, the number of POIs in the gas station's neighborhood, and (2) *distance to POI*, the minimum average distance between a candidate's points and nearby POIs.

We use a manually labeled dataset to train a gradient tree boosting classifier [Friedman 2002] and then use the trained model to distinguish real REs from other behavior.

6. EXPECTED DURATION LEARNING

A unit's expected duration is an indicator that shows the average time spent at a gas station during a certain period. Currently, we have discovered taxi REs from the raw trajectory dataset. If there are enough REs incorporated, we can use their average

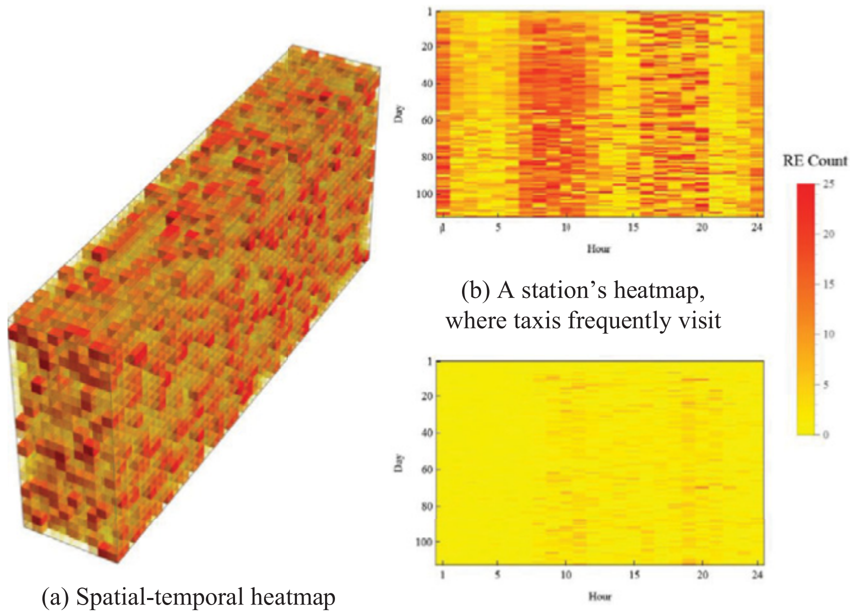

(a) Spatial-temporal heatmap

(b) A station's heatmap, where taxis frequently visit

(c) Another station's heatmap, where taxis rarely visit

Fig. 8. Detected REs' heatmap in spatial-temporal space.

values to estimate this indicator. However, only a portion of units are filled with enough detected REs, considering that there are thousands of gas stations and the taxis' finite REs cannot cover each gas station at each moment. Figure 8 reveals the exact scenario for the entire space as well as two slices along the gas station dimension, where the color signifies the number of detected REs. Even though the gas station in Figure 8(b) is popular for taxis, during some periods there are still few taxis. The situation is even worse for the station in Figure 8(c). To predict a less occupied or even empty unit's expected duration, we apply a context-aware collaborative filtering model to solve the data sparsity problem and then detail how to improve the performance with the aid of various contextual features.

6.1. The CATF Model (Context-Aware Tensor Factorization)

Currently, for units with enough detected REs, their expected durations are obtained and treated as observable data. Our concern is then to find the remaining units' expected durations, which is similar to predicting the missing ratings in recommender systems [Ricci and Shapira 2011]. Thus, we utilize the factorization method, which is widely adopted in dealing with user-item rating problems. Since there are three dimensions in our space, we first discuss how to apply tensor factorization to predict a unit's expected duration.

We apply high-order singular value decomposition (HOSVD) [De Lathauwer et al. 2000] to factorize the three-dimensional tensor into three matrices, $H \in R^{h \times d_H}$, $G \in R^{g \times d_G}$, $D \in R^{d \times d_D}$, and one central tensor, $S \in^{d_H \times d_G \times d_D}$, where d_H, d_G, d_D are dimensionality parameters to balance capability and generalization. The reconstructed value for unit U_{ijk} in the traditional tensor factorization [Karatzoglou et al. 2010] is given as

$$F_{ijk} = S \times_H H_{i*} \times_G G_{j*} \times_D D_{k*}. \tag{1}$$

We denote the tensor matrix multiplication as $\times_M$, where the subscript denotes the direction (i.e., $T = Y \times_M M$ is $T_{ljk} = \sum_{i=1}^{h} Y_{ijk} \times M_{ij}$). The entries of the ith row of matrix M are represented as M_{i*}.

The factorization tends to drill out the hidden relationship based on the observations. For instance, if two gas stations' observed expected durations are alike on most days at a certain hour, the factorization method will guide them to perform similarly on the remaining days at that hour, where the similarity is measured on the latent subspace controlled by dimensionality parameters d_H, d_G, d_G. Additionally, the single tensor factorization does not take full advantage of our data, as it only relies on the already observed expected durations to fill up the blank ones. Another important signal, what we call *contextual features*, such as gas stations' attributes and weather conditions, which are the essential factors that determine a user's refueling choice, has not been considered. Actually, economists have found that the stations' exogenous and endogenous factors (location, nearby traffic flow, size, brand, price, etc.) have a great effect on gas stations' competitive conditions [Iyer and Seetharaman 2005], and thus these factors could indirectly influence the time spent at the gas stations (if a gas station is popular, a longer wait time might be needed). In addition, workdays versus holidays, price-fluctuating days versus normal days, and even the weather conditions can have an impact on a person's refueling decision and affect the time spent at the gas stations.

The user's or item's contextual features are often modeled in recommender systems to help reduce uncertainty issues [Ricci and Shapira 2011]. We apply a similar method to integrate the contextual features with factorization in our CATF model. Assume that there are c_1, c_2, c_L features, where feature c_l has categorical values $1, 2, z_l$ to refer to possible contextual conditions. Therefore, integrating the tensor factorization with the context features [Baltrunas et al. 2011], the reconstructed value in CATF for unit U_{ijk} is redefined as

$$F_{ijk} = S \times_H H_{i*} \times_G G_{j*} \times_D D_{k*} + \sum_{l=1}^{L} B_{lc_l}, \tag{2}$$

where B_{lc_l} is the parameter modeling how this unit's contextual feature l with condition c_l would have an effect on the reconstructed value. The introduced contextual parameters guarantee the fact that units with similar contextual features tend to have similar time spent (the part $\sum_{l=1}^{L} B_{lc_l}$ tends to be similar between similar units). To generate time spent predictions, the model parameters should be learned using the observable data. We define the learning procedure as an optimization problem:

$$\min_{H,G,D,S,B} L(Y, F) + \Omega(H, G, D, S, B), \tag{3}$$

where $L(Y, F)$ is the loss function given as

$$L(Y, F) = \frac{1}{||S||_1} \sum_{i,j,k} Z_{ijk} \cdot (Y_{ijk} - F_{ijk})^2, \tag{4}$$

where $Z \in \{0, 1\}^{h\times g\times d}$ is a binary tensor with nonzero entries Z_{ijk} whenever Y_{ijk} is observed. Equation (4) indicates that we should consider the reconstructed accuracy for observed units. $\Omega(H, G, D, S, B)$ is the regularization term to prevent overfitting, which is given as

$$\Omega(H, G, D, S, B) = \frac{1}{2}\lambda \times \left(||H||^2_{Frob} + ||G||^2_{Frob} + ||D||^2_{Frob} + ||S||^2_{Frob} + ||B||^2_{Frob}\right). \tag{5}$$

Regularization is controlled by the meta parameter λ. We use grid search to find model parameters (note that we set $d_H = d_G = d_D$ in our experiment setup) and stochastic gradient descent [Zhang 2004] to solve this optimization problem.

6.2. Contextual Features Engineering

Various factors, which are related to the refueling decision and thus can influence the time spent, are explored in detail here. From the standpoint of stations, five types of contextual features, including POI, traffic, brand, price, and the size of the area, are taken into account. Beyond that, we also consider three types of time-dependent contextual features: price-fluctuated days versus normal days, workdays versus holidays, and weather conditions. We elaborate on how to construct these features and reveal the influence that they exert on the time spent.

Point-of-interest feature. According to Monod et al. [2011], geographic concentration and co-location of industries will significantly affect the economic activities. Analogously, drivers' refueling decisions can be shaped by the gas stations' surrounding environments. We determine this effect based on a gas station's nearby POIs. For each category C of the POIs, to discover its correlation to the gas station, we use the metrics defined in Jensen [2006], which is given as

$$J_c = \frac{\#co\_location(C, g)}{\#C}, \tag{6}$$

where $\#co\_location(C, g)$ refers to the frequency of co-location for category C with the gas station and $\#C$ indicates the individual frequency. The top five discovered POIs are {Service Zone At Motorway, Toll station, Factory, Vehicle Maintenance, and Vehicle Service}.

Aggregating nearby POIs, the POI feature of a gas station is given as

$$F_p(g_i) = \sum_C N(C, g_i) \cdot J_C, \tag{7}$$

where $N(C, g_i)$ indicates the frequency of the category C standing by station g_i.

Traffic feature. The traffic feature of a gas station depends on its nearby traffic flow and competitive conditions. By aggregating all trajectory data for each road, we can estimate a road's traffic flow. We determine how a road's traffic flow influences nearby gas stations based on the Huff probability model [Huff 1964], which is given as

$$TF(r \to g_i) = TF_r \cdot \frac{\frac{1}{dist(g_i, r)}}{\sum_{g_j} \frac{1}{dist(g_j, r)}}. \tag{8}$$

Finally, the traffic feature of a gas station is given as

$$F_T(g_i) = \sum_r TF(r \to g_i). \tag{9}$$

Brand feature. Having a strong brand position means that the brand has a unique, credible, sustainable, and valued place in the customer's mind [Davis 2000]. It has been disclosed that India's petroleum market has been dominated by state firms [Attri et al. 2011]. In addition, a survey of 218 Polish consumers has been conducted, and the findings suggest that consumer ethnocentrism is displayed by a more positive perception of the domestic brand even in a situation where foreign brands are superior to domestic ones [Supphellen and Rittenburg 2001]. The situation is similar in

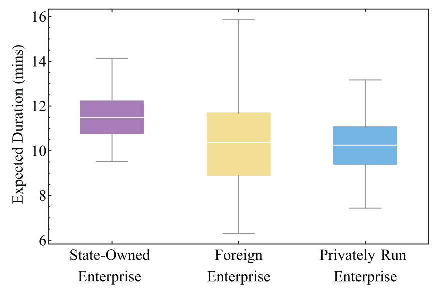

Fig. 9. Brand effect for gas station.

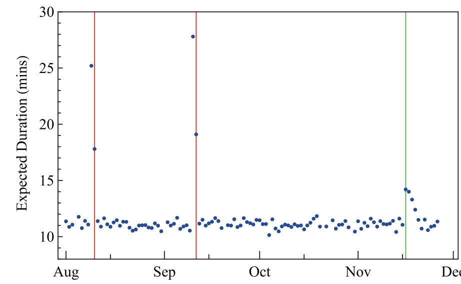

Fig. 10. The effect of price fluctuation.

China. Two state-owned enterprises, Sinopec[2] and CNPC,[3] have been dominating the domestic petroleum market and exert a powerful brand effect on customers. We extract the stations' brands from their name description and put them into three classes: State Owned, Foreign, and Privately Run. Figure 9 shows the quantile distribution of time spent by taxi drivers for each class, which indicates that state-owned enterprises display higher time spent than the other two classes. Taxi drivers have also revealed that they prefer the state-owned companies because the petroleum quality is more trustworthy.

Price feature. Price is a sensitive factor affecting personal consumption choices. Chan et al. [2007] showed that in Singapore, fuel consumers are willing to travel up to a mile for savings of \$.03 per liter. However, petroleum service stations rarely publish the price conditions due to commercial considerations. We collected the price data from the self-reports dataset. As shown in Figure 4(b), the price is calculated according to the division between refueling fees and refueling charges. To eliminate the inconsistent prices at a station, which are aroused by the discrepancy in prices among different petroleum products and the intense price fluctuation (caused by price regulations, which will be detailed later), we only focus on the prices of the most widely consumed 93# petroleum and select those filled in from September 12 to November 15 (no regulated price adjustment during this period). Finally, we use the average value collected at a gas station to indicate this station's price feature.

Area feature. The area feature of a gas station reflects its passenger capability, which influences the time spent at this station. We manually labeled the gas stations' areas by using the distance measurement tool on satellite maps.

Price fluctuation feature. Price regulations on the petroleum market are a distinct policy with strong Chinese characteristics. Since 1998, China has regulated oil product prices to maintain inflationary pressures, maintain a demand and supply balance, and smooth out price volatilities from international markets [Shaofeng 2006; Ling-Yun and Yan 2009]. Although different gas stations can provide different prices, they are forced to rise or cut prices simultaneously to respond to the government's requirements at a particular moment, which will exert intensive influence on customers' refueling decisions. According to the price adjustment record in Beijing,[4] Figure 10 how price shock affected taxi drivers' refueling time spent with regard to the detected REs in our data over a 4-month period in 2012 (a red line refers to a day with a price increase, and a green line refers to a day with a price decrease). The two moments with price increases, August 10 and September 11, both caused the durations to rise suddenly

[2]China Petrochemical Corporation: http://www.sinopec.com.
[3]China National Petroleum Corporation: http://www.cnpc.com.cn/en.
[4]http://data.eastmoney.com/cjsj/yjtz/beijing.html.

Table I. Analysis of Variance (ANOVA) for Testing the Correlation between Contextual Features and the Expected Duration

Feature	POI	Traffic	Brand	Price	Area	Weather	Price Fluctuation	Workday vs. Holiday
F-value	367	862	442	475	184	79	154630	308
P-value	$<e^{-16}$	$<e^{-16}$	$<e^{-16}$	$<e^{-16}$	$<e^{-16}$	$<e^{-16}$	$<e^{-16}$	$<e^{-16}$

and sharply. In particular, the figure shows that the day before the price increase is much more intensive than the day of the price increase itself. This phenomenon is interpreted by the fact that the government announces price increases one day before they actually take effect, and thus customers flood into stations to save money following the announcement.[5] Interestingly, a decrease displays a different pattern. Figure 10 reveals that from November 16 to November 19, due to a price decrease, more drivers chose to refuel in that period and the time spent only rose a little more than usual. In this situation, customers did not hurry to stations before the announcement came into force because they could wait and save money. Moreover, the effect of a price decrease is smoother and lasts longer, so it can be reasoned that this situation is not as urgent as a price increase and customers can enjoy the tangible benefit for a long time. In a nutshell, Figure 10 indicates the effects of price fluctuations on refueling times. To incorporate the price fluctuation feature, we classify the timestamp of the day into three categories: normal days; days influenced by price increases, including the day and the day before; and days influenced by price decreases, including the day and the following 3 days.

Other features. Out of consideration for the effect of weekends and festivals, we label the timestamp of a day as a workday or holiday. Furthermore, to take into account the influence of weather conditions, we also crawled the weather history Web site[6] to classify each day as rainy or nonrainy.

After the engineering procedure, considering that CATF needs categorical variables, we divide the continuous features (POI, traffic, price, and area) into three categories separately as the final adopted contextual features (notice that the seven features correspond to $c_1, \ldots, c_8$ separately). Table I details the statistical correlation between each feature and the time spent, showing that all of these features exert a prominent effect. In particular, for the price fluctuation feature, the supranormal F-value accounts for its drastic impact on customers' refueling decisions.

7. ARRIVAL RATE CALCULATION

A unit's expected duration indicates the time spent there. We also want to know how many vehicles have visited this unit, from which we can estimate the energy consumption. However, our dataset only covers about 30,000 taxicabs, which is just a small portion of the total number of vehicles in the city. To solve the sparsity problem inside a gas station, we estimate the total arrival rate by modeling each gas station as a queue system.

7.1. Queue System

Generally speaking, there are several queues at a station, and each queue could simultaneously serve several vehicles. To reduce the complexity of the system, we make some simplifications. First, we ignore transfers from one queue to another queue and assume that each vehicle is fixed to a certain queue. This assumption guarantees that each queue can be treated as an independent queue system. Moreover, we make the

[5]http://auto.msn.com.cn/pic/1364298.shtml.
[6]http://www.wunderground.com/history.

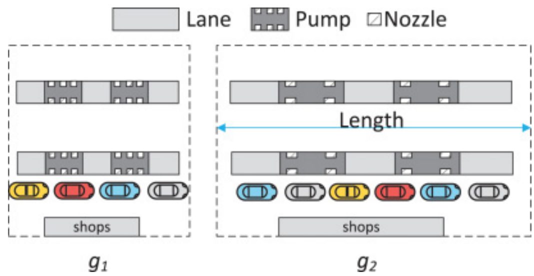

Fig. 11. Layout of gas stations.

assumption that drivers will always choose the shortest queue to join, which ensures that each queue will share the same wait time on the whole.

Assume that there are Q queues at the gas station. In a particular unit related to this station, the vehicles' arrival flow for each queue Q_i is described as a homogeneous Poisson process $N(t, \lambda_i)$, which indicates that the number of vehicles in the period $[0, t]$ is a Poisson distribution with parameter $\lambda_i \cdot t$ [Kleinrock and Wiley 1977]. The unit of t is hours, which is the same as the period of a unit. Thus, λ_i is the number of vehicles that joined this queue in this unit, and the overall arrival rate of this unit is given as $\lambda = \sum_{i=1}^{Q} \lambda_i$.

7.2. Calculation

In the queue system, given customers' arrival stochastic process and servers' service time distribution, the equilibrium indicators such as wait time and system time can be obtained [Kleinrock and Wiley 1977]. We assume that all refueling equipment is undifferentiated and that its service time satisfies exponential distribution $Exp(\mu)$. For the ith queue Q_i, we assume that it has c_i servers, and thus this queue can be treated as an $M/M/c_i$ system. Its average arrival rate is λ_i, and its average service time is $\frac{1}{\mu}$. The equilibrium indicators can be computed as follows [Kleinrock and Wiley 1977]:

$$W_s = \frac{\frac{\lambda_i^{c_i}}{\mu^{c_i+1}(c_i-1)!}}{(1-\frac{\lambda_i}{c_i\mu})^2} \cdot \left[\sum_{k=0}^{c_i-1} \frac{1}{k!} \cdot \left(\frac{\lambda_i}{\mu}\right)^k + \frac{1}{c_i!} \cdot \frac{1}{1-\frac{\lambda_i}{c_i\mu}} \cdot \left(\frac{\lambda_i}{\mu}\right)^{c_i}\right]^{-1} + \frac{1}{\mu}, \tag{10}$$

where W_s is the equilibrium system time (including both the wait time and service time), which indicates the amount of time it needs to stay there during equilibrium. Since we believe that drivers are rational, each queue's equilibrium system time is the same, and we use this unit's estimated expected duration to represent W_s. It is obvious that W_s depends on λ_i, μ, and c_i. Therefore, given μ, c_i, and W_s, we solve the equation to get parameter λ_i, and the equation can be solved by a numerical algorithm, such as the Newton Raphson method. Finally, this unit's arrival rate λ is gathered by each queue's corresponding λ_i.

7.3. Parameters Determination

We assume that the shortest duration of all detected REs corresponds to the service time (there are some cases where taxis can refuel directly). We select the top 1,000 shortest durations and use their average value to estimate $\frac{1}{\mu}$. We then need to determine Q (number of queues) and c_i (number of servers in a queue) at each gas station, which depends on the gas station's area, the arrangement of pumps, and how many nozzles, as shown in Figure 11. A pump has several nozzles, and the nozzle plays the role of

Table II. Temporal Distance between Candidate and Real RE

Temporal Distance (minute)	HLD-1		HLD-3	
	Mean	Std.	Mean	Std.
$\lvert r.AT - c.AT \rvert$	1.05	0.42	0.49	0.31
$\lvert r.DT - c.DT \rvert$	1.24	0.51	0.70	0.21
$\lvert r.AT - c.AT \rvert + \lvert r.DT - c.DT \rvert$	2.29	0.47	1.19	0.25

server. As mentioned earlier, we measure the stations' lengths in satellite maps. We also go through the street view maps to observe the number of lanes N_l and the number of nozzles along a queue N_{n_i}. We see that a pump can serve both sides simultaneously, and therefore Q is equal to $2 \times N_l$. It is a little tricky to determine c_i. The figure shows that gas station g_1 has six nozzles along a queue. However, due to the length limitation, it can only serve four vehicles simultaneously. The situation is the opposite for g_2. Thus, we set $c_i = Min(N_{n_i}, \frac{length\ of\ gas\ station}{length\ of\ a\ vehicle})$. In reality, the length of a normal automobile is about 4.5, and therefore we set it at 5m in view of the gap between vehicles.

8. EXPERIMENT

In this section, we first describe the human-labeled datasets and then evaluate the performance of RE detection, the expected duration learning, and the arrival rate calculation.

8.1. Human-Labeled Dataset Description

We employ four human-labeled datasets for both a training model and evaluation as follows.

(1) *HLD-1*: We manually labeled 1,500 real REs by plotting the taxis' raw trajectories on digital maps. In consideration of possible bias, these labeled REs are randomly selected from trajectories including different regions and time. In the end, 1,200 of them were used to learn the parameters in candidate extraction, and the remaining were used to validate the performance.

(2) *HLD-2*: We manually labeled 2,000 candidates (True/False) by plotting the extracted candidates on digital maps.

(3) *HLD-3*: This dataset covers 33 trajectories collected by two authors, and each trajectory contains a recorded RE (arrival time, departure time, selected gas station), which is used to evaluate the performance of RE detection.

(4) *HLD-4*: To evaluate whether the expected duration learning component and the arrival rate calculation component work well in reality, we chose two gas stations on which to perform a field study. We recorded the vehicles' arrival and departure times (there were many vehicles, and we could not record all of their information, so we just selectively recorded some cases), as well as how many vehicles had refueled there in that period. This field study lasted from October 17 to November 15 in 2012, ranging from 5:00 pm to 6:00 pm each time. In total, 14 days of records were collected (each station had 7 days' worth of records).

8.2. Experiments for Refueling Event Detection

In this section, we evaluate the effectiveness of candidate extraction and the filtering model separately.

8.2.1. Results of Candidate Extraction. We used 1,200 instances in HLD-1 to set the parameters by grid search and evaluated the performance both on the remaining 300 instances in HLD-1 and the authors' collected dataset HLD-3. As shown in Table II, we computed the temporal distance (AT corresponds to arrival time and DT corresponds to departure time) between the labeled refueling time and the nearest candidates

Table III. Results of Filtering

	Features	Precision	Recall
HLD-2	Nonfiltering	0.465	1.0
	Spatial	0.626	0.729
	Spatial+Temporal	0.893	0.858
	Spatial+Temporal+POIs	0.916	0.906
HLD-3	Nonfiltering	0.825	1.0
	Spatial	0.875	0.848
	Spatial+Temporal	0.941	0.969
	Spatial+Temporal+POIs	0.941	0.969

Table IV. Number of Detected REs in Each Unit

	D_1	D_2	D_3	D_4	D_5	D_6	D_7
g_1	7	6	5	5	6	6	4
g_2	0	1	0	0	0	0	2

discovered. The performance was better in HLD-3 because the GPS devices used by the authors had a lower sampling interval (the sampling interval was about 5 seconds, whereas the taxis' GPS sampling interval was about 1 minute).

8.2.2. Results of Filtering. The precision and recall with respect to features that we used for the classifier are presented in Table III. We applied a 10-fold cross-validation method on dataset HLD-2. The performance on HLD-3 was still better than HLD-2 because there was less noise in the candidates. Compared to private car owners, taxi drivers visit the POIs near to gas stations more frequently, including vehicle repair shops or parking lots, which can generate pseudocandidates. As well, we found that the temporal feature plays an important role in both datasets. In any case, the precision and recall were both higher than 90%, which is accurate enough for the next step. After applying the method to all candidates, we extracted 1,342,957 REs in total. The statistics of the results shows that a taxi driver refueled about every 2 days, the average time spent was 10.57 minutes, the minimum was 3.51 minutes, and the maximum was 60.18 minutes.

8.3. Experiments for Expected Duration Learning

There were a total of 1,852,032 units (24 hours $\times$ 689 gas stations $\times$ 112 days), and each unit incorporated about 0.725 REs on average, which indicates that many units lacked enough detected REs to estimate the expected duration. Table IV details how many detected REs were covered during the period of our case study at these two gas stations. g_1 was more attractive to taxis, and these units incorporated enough detected REs, whereas taxis rarely patronize g_2. Therefore, for each unit in g_1, its expected duration is represented by the detected REs' average duration, and the results are shown in Figure 12(a) and compared with the results of the recorded vehicles' duration in g_1 in the field study. The standard deviation of records is about 2 minutes, which shows that over the course of an hour, the refueling time spent is almost stable.

We used four baselines for the comparison:

Average filling (including AWH, AWD, and AWG). (1) AWH (*average within hour*): For a unit without sufficient detected REs, AWH finds all other units with the same hour timestamp and uses their average expected durations to estimate this unit's expected duration. (2) AWD (*average within day*): Similar to AWH, AWD uses the average value within the same day. (3) AWG (*average within gas station*): Analogous to the previous two methods, AWG uses the average value within the same gas station.

Support vector machine (SVM). SVM uses the seven previously represented contextual features of a unit, as well as the timestamp of the hour and the timestamp of

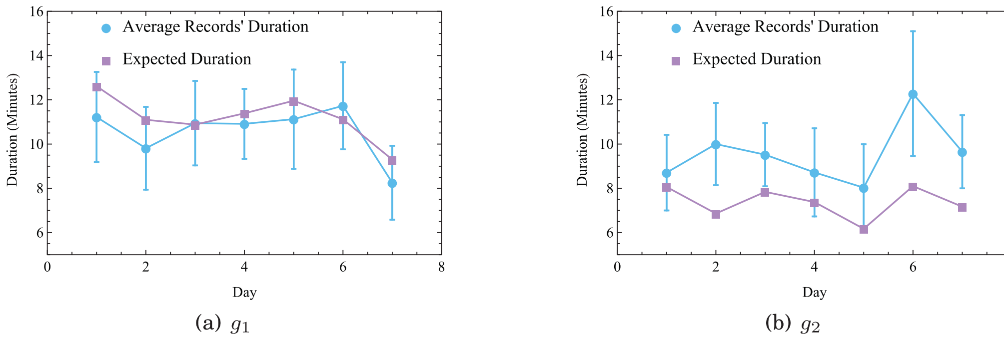

Fig. 12. Records' duration with respect to expected duration.

Table V. Results of CATF with Respect to Baselines

	MeanErr	Std.
AWH	3.02	0.97
AWD	3.73	1.27
AWG	3.11	1.14
SVM	3.08	1.06
TF	2.66	0.82
CATF_1 (TF + POI)	2.48	1.02
CATF_2 (TF + POI + Traffic)	2.27	0.87
CATF_3 (TF + POI + Traffic + Brand)	2.24	0.91
CATF_4 (TF + POI + Traffic + Brand + Price)	1.93	0.87
CATF_5 (TF + POI + Traffic + Brand + Price + Area)	1.88	0.84
CATF_6 (TF + POI + Traffic + Brand + Price + Area + Price_fluctuation)	1.86	0.94
CATF_7 (TF + POI + Traffic + Brand + Price + Area + Price_fluctuation + Others)	1.83	0.81

Table VI. Price Fluctuation's Effect on CATF Results

		MeanErr	Std.
Normal Days	CATF	1.82	0.83
	CATF Except Price_fluctuation	1.83	0.85
Price Increase Influencing Days	CATF	2.77	1.86
	CATF Except Price_fluctuation	4.55	1.35
Price Decrease Influencing Days	CATF	2.17	1.09
	CATF Except Price_fluctuation	3.13	1.17

the day as temporal features to train a supervised model using SVM regression. We selected the units that incorporate more than two detected REs and finally obtained 651,178 units as observable data. We evaluated our model using 10-fold cross validation on the observable data and used grid search to set $d_H = d_G = d_D = 13$ and $\lambda = 0.8$. The results are presented in Table V, where MeanErr signifies the average offset between the observable value and the predicted value for all testing data in the 10-fold cross validation. The unit of MeanErr is minute and similar to Std. The table shows that the contextual features play an important role in improving the performance, especially the POI, traffic, and price features. It seems that the price fluctuation factor had little effect on the results, because its impact is transitory and the effect on the results is smoothed by the normal days, which have numerical superiority. Table VI details the price fluctuation effect on different types of days, which reflects the price fluctuation of the influencing days. The time spent was much greater and varied more heavily than usual. Table VI also clearly presents the importance of the price fluctuation feature

Table VII. Description of Two Gas Stations

	N_L	N_{n_i}	*Length*	Q	c_i
g_1	3	4	27.2m	6	4
g_2	2	4	18.7m	4	3

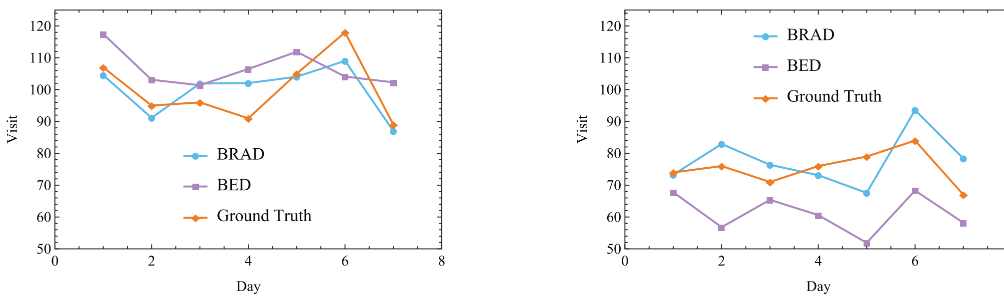

(a) g_1 (b) g_2

Fig. 13. Results of arrival rate.

during the days influenced by the price fluctuations. As shown in Table V, the SVM model performed even worse than AWH, perhaps because the data tensor is quite sparse and a supervised model is not suitable for this situation. The results indicate that the error in time spent estimated in a unit could be limited to about 2 minutes on average.

Additionally, to evaluate the performance of CATF in reality, we compared the predicted value with the records of g_2 in Figure 12(b). It seems that the error was a little larger than estimating the expected duration directly in Figure 12(a).

8.4. Experiments for Arrival Rate Calculation

In this section, we discuss the calculation of the units' arrival rate. Table VII details the records of two gas stations, as well as their determined queue-model parameters. These two gas stations have an identical number of nozzles in each queue, denoted as N_{n_i}. Similarly, each queue's number of servers is denoted as c_i. For the service time parameter μ, we selected the top 1,000 shortest durations among all detected REs and obtained $\mu = 4.01$ minutes.

We compared the following methods with the ground truth (the total recorded vehicle visits at two gas stations for each day):

—*BRAD* (*based on recorded average duration*): This method uses the selectively recorded vehicles' average duration to estimate equilibrium system time W_s.
—*BED* (*based on expected duration*): This method makes use of each unit's expected duration to estimate W_s.

The results are shown in Figure 13. The figure shows that BRAD is approximate to the ground truth, which illustrates the effectiveness of our queue system model. In addition, the figure indicates that BRAD is more accurate than BED, because BED depends on the results of RE detection and expected duration learning, and the errors accumulated in these two parts exert an influence on the results of the arrival rate. However, for both gas stations, we found that the gap between BED and the ground truth was acceptable.

9. URBAN REFUELING ANALYSIS

So far, we have detected the REs of taxi drivers (in our dataset), collected tens of thousands of online users' (representing a small-scale sampling of household consumers)

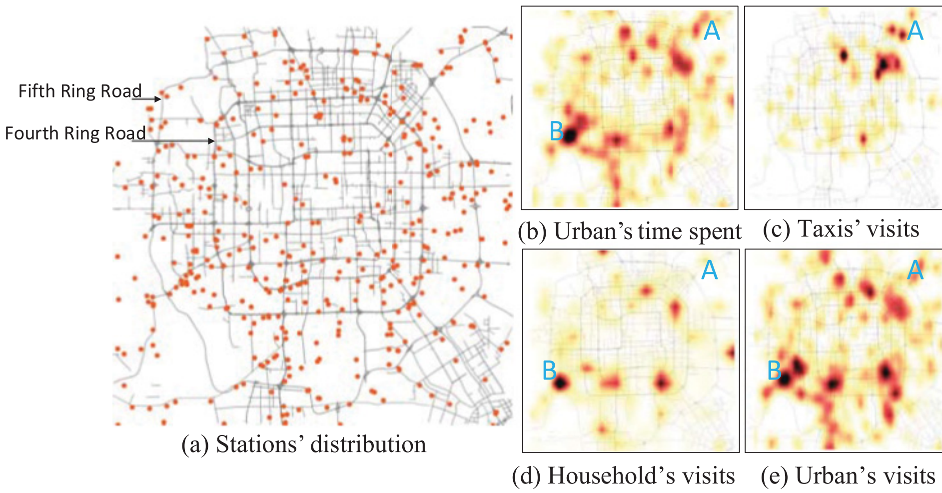

(a) Stations' distribution (b) Urban's time spent (c) Taxis' visits (d) Household's visits (e) Urban's visits

Fig. 14. Refueling behavior's spatial distribution.

refueling consumption, and obtained two indicators—expected duration and arrival rate—for the city. We can determine the refueling behavior of different groups from a spatial, temporal, and economic view.

9.1. Geographic View

Figure 14(a) reflects how gas stations are scattered throughout the city. The gray lines depict the city's road network. The figure shows that a large portion of stations are located between the fourth ring road and the fifth ring road, whereas fewer stations are distributed in the central part of Beijing.

For each gas station, we use the average expected duration and visits to denote the time spent and visit separately. Since the estimation of urban time spent is based on the taxi drivers' time spent, the spatial distribution of time spent for taxis is similar to that for the whole city and thus we only plot the urban time spent distribution as shown in Figure 14(b). The three other figures show the spatial distribution of visits for taxis, household customers, and the whole city separately. A darker color refers to a longer time spent or more visits. Comparing Figure 14(c) to Figure 14(b), it shows that most of the areas that taxi drivers frequently visited also had a longer time spent there. On the other hand, taxis drivers rarely patronized stations in area B. However, a long wait time was still required, which implies that there were many other vehicles refueling there, which is verified by our household customer data shown in Figure 14(d). Figure 14(c) and Figure 14(d) show that taxi drivers and household customers have different refueling patterns in terms of the spatial aspect. According to our survey, we found that area B was near the entrance to a major highway, and thus many household customers refueled there. Taxi drivers frequently refueled on the southeast part of the fifth ring road or several other small-scale hotspots scattered to the south and north. Actually, these hot areas are transportation hubs and are more likely to attract taxis. For instance, the popular area A is near the highway directly to the airport, where many taxis travel, and thus there is a higher probability that they will refuel at nearby gas stations. Figure 14(d) shows that the hotspots of the whole city can cover the places frequently visited by taxi drivers and household customers, such as area A and area B. Furthermore, comparing Figure 14(b) to Figure 14(a), it reveals that although a large

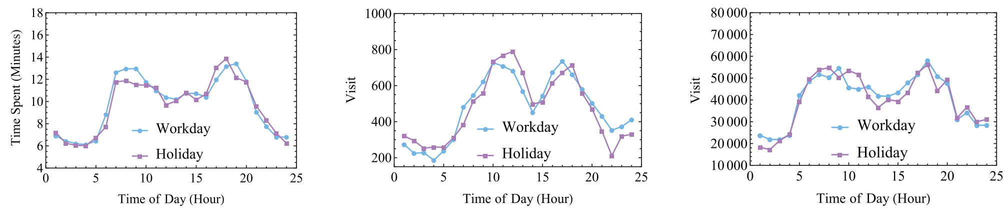

(a) Urban's time spent (b) Taxis' visits (c) Urban's visits

Fig. 15. Refueling behavior's temporal distribution.

number of stations have been built in area B, the long time spent suggests that new stations still should be planned nearby.

9.2. Temporal View

We compared the average time spent and visits at different hours of the day and considered workdays and holidays separately. Since taxis' time spent distribution is similar to that of the whole city, we only plotted the urban distribution as shown in Figure 15(a). On the one hand, during rush hour (7 am, 8 am, 6 pm, 7 pm), many cars came to refuel and a longer wait time was required. On the other hand, the figures show that on holidays, a shorter wait time was needed at 7 am and 8 am, whereas a little more time was needed at 9 am and 10 am. This phenomenon accords with the habits of office workers, as they often choose to refuel on the way to work in the morning. They do not need to wake up early on holidays, and therefore there were fewer early customers. Considering that for the refueling reports of household customers the granularity of the recorded time is days rather than minutes or seconds, we only compare the visit distribution of taxis and the whole city. Figure 15(b) shows the refueling climax of taxi drivers was about 10 am, which indicates they chose to stagger the busy period at about 8 am. The two peaks in Figure 15(c) indicate higher petrol consumption during these periods and warn people to avoid refueling at that time.

9.3. Economic View

In keeping with the aforementioned price feature, we also focus on the price of 93# petroleum from September 12 to November 15 (the prices are partitioned into three categories). In Figure 16, the gray bars show the distribution of gas stations with the three-categorized price. For each satisfactory refueling report of household customers, the direct labeled refueling price was categorized to a certain class and their refueling distribution is shown in the yellow bars. For taxi drivers, since we are aware of the place where they refuel and the price feature related to each station, their refueling distribution is shown in the red bars of Figure 16. The figure shows that taxi drivers prefer low-price consumption, indicating that compared to normal private car owners, the taxi drivers' financial situation is tighter. From another point of view, most refueling consumers, whether taxi drivers or not, would like to choose the middle-price stations, maybe because the drivers do not trust the quality of the petroleum provided by low-price stations. In addition, the quality provided by the middle-price stations is good enough, and they do not need to seek the expensive product. For the entire city's results, we first obtained each station's visits from our model and mapped the stations' prices to obtain the distribution. Since the urban result reflects the whole city's refueling pattern, the figure is more similar to the household customers' preferences. This may be because the number of household customers dominates the fuel market.

Figure 17 delineates how the price fluctuation factor effects different groups' refueling decisions. It shows that taxi drivers are hardly influenced by this factor, as they

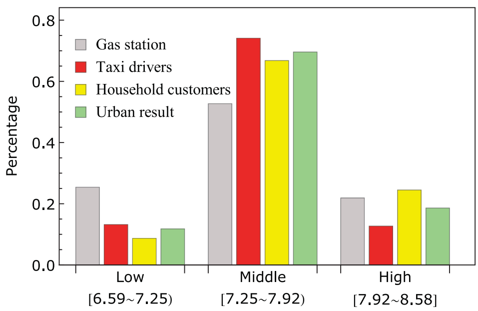

Fig. 16. Refueling distribution of three-categorized price. The unit of price is Yuan per liter.

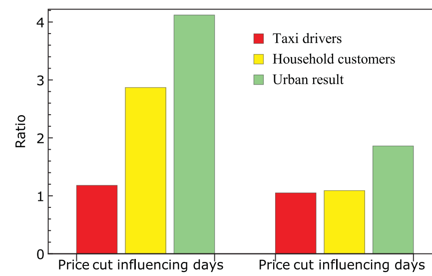

Fig. 17. Refueling frequency influenced by the price fluctuation factor. The ratio value in the vertical axis indicates the average visit on price fluctuation days divided by the average visit on normal days.

actually have to refuel every other day. However, the situation is different for household customers, who refuel about twice a month. Thus, they would like to refuel in advance or delay more than usual to save money when faced with price fluctuations. The overall results are even more drastic than those of household customers. This may be because the user we extracted from LiCheWang is only a small subset and does not reflect the entire city.

10. DISCUSSION

In this section, we discuss the general application of our methods and the limitation of the system.

Our method was currently evaluated on taxis' trajectories. However, other vehicles can seamlessly be incorporated into this system. During the refueling detection phase, as shown in the experiments, we found that the results for private cars outperform that for taxis, because taxi drivers tend to engage in other activities near gas stations. When we obtained the two indicators in a unit, we actually only relied on the detected REs' time spent, which is the result of RE detection and is independent of whether this vehicle was a taxi or not. Our taxi trajectories can be regarded as a sampling of the whole trajectories generated by all vehicles in the city.

Additionally, taxi drivers form a special group for refueling consumption. They have biased preference, as shown in the urban refueling analysis section, which will lead them to some units while aggravating the sparsity issues of other units. The potential difference in refueling regularity between taxi drivers and normal drivers might degrade the accuracy. We also use taxis' refueling time solely to estimate the parameter of refueling time distribution, which might bring some bias (some other vehicles' refueling time will usually be longer than that of taxis, such as trucks). In addition, the drivers' behavior in the gas stations' queue system is ideally assumed, and the reality is usually more complex than we can capture. For instance, when the lane in a gas station is narrow, a car that has finished refueling might be blocked by the car in front, and this special case is difficult to incorporate into our system.

11. RELATED WORK

11.1. Customer Refueling Behavior Analysis

A comprehensive picture of refueling behavior analysis provides insight for fueling marketing strategy and government planning. However, as pointed out by Memmott [1963], refueling behavior itself has been constantly neglected in previous empirical and theoretical studies of travel behavior and fuel use. In the limited refueling behavior research, survey-based methods have been applied to focus on understanding customers' refueling habits to help make decisions. Kitamura and Sperling [1987], which is the earliest work that we can track, presented a descriptive and explanatory analysis of refueling behavior of gasoline-car drivers in northern California, and the results revealed that refueling behavior is only weakly correlated with demographic and socioeconomic descriptors of drivers. Kelley and Kuby [2013] interviewed 259 drivers in southern California and analyzed their refueling behaviors. The results were used to help select appropriate optimal facility location models. Li et al. [2012] used a smartphone application to build a driving behavior monitoring and analysis system for hybrid vehicles. Compared to their interesting and influential work, which primarily was aimed at individuals, our system steps further toward macroscale analysis through large-scale datasets.

11.2. Gas Station Analysis and Planning

Gas station problems mostly concentrate on facility location, which is to investigate how to build a new gas station or analyze the current gas station layout in consideration of various factors, such as price, geographical distance, and the surrounding environment. For instance, Chan et al. [2007] proposed an econometric model to analyze both the geographic locations of gasoline retailers in Singapore, as well as price competition between these retailers. Their results revealed that in Singapore, fuel consumers are willing to travel up to a mile for savings of $.03 per liter. Iyer and Seetharaman [2005, 2008] examined how product design, prices, and location characteristics influenced price competition in retail gasoline markets, which uncovered the various factors that provide insight for the gas station planning problem. In addition, there has also been an extensive body of research on electrical vehicle–related issues, which have looked at where to place the fueling stations and dynamics of a user's mobility across the city. Hatton et al. [2009] reflect the essential role of supportive infrastructure in the mass implementation of electric vehicle technology, which is to provide efficient and diverse recharging solutions for vehicle drivers. The authors discussed the related problem areas, ongoing developments, and future trends in the design and development of charging systems. McPherson et al. [2011] described a model for identifying the optimal geographic locations for charging stations and how best to stage the rollout over an extended time period. The previous works mainly concentrated on analyzing stations' attributes, whereas our work tries to discover stations' petrol consumption through passive human sensing in a more intuitive way to understand the stations' operating status.

11.3. Urban Computing

With the popularity of diverse sensors, exploring the rule of city is a burgeoning and attractive area in computer science. The term *urban computing* has emerged to concentrate on the integration of computing, sensing, and actuation technologies into everyday urban settings and lifestyles. In recent years, on the one hand, a great deal of work based on spatial-temporal analysis has been proposed to explore the status of the city [Zheng et al. 2011; Kindberg et al. 2007]. For example, the auto-GPS mobile sensor data of approximately 1.6 million people throughout Japan were applied to discover the short- and long-term evacuation behaviors for individuals [Song et al. 2013]. In one paper by Ge et al. [2011b], taxi GPS traces were utilized to effectively uncover taxi drivers' fraud activities; in another paper, Ge et al. [2011a] developed a taxi business intelligence system to explore massive taxi location traces from different business perspectives with various data mining functions. On the other hand, the knowledge of the urban exploration has been applied for providing various services to individuals. For example, Yuan et al. [2010] provided a strategy of finding efficient driving directions based on the knowledge of taxis drivers, and a more recent work [Yuan et al.

2012] presented a recommender system for both taxi drivers and passengers based on passengers' mobility patterns and taxis drivers' pick-up/drop-off behavior. Yang et al. [2013] took advantage of crowdsourced digital footprints and ratings to subtly characterize an individual's fine-grained location preference; the amount of online travel information was also collected for personalized travel package recommendation in the work of Liu et al. [2011]. Our work concentrates on catching a glimpse of urban transportation's energy consumption, which is a concern of urban computing.

12. CONCLUSION AND FUTURE WORK

In this article, we propose an intelligent system for discovering urban refueling behavior using taxis trajectories, POIs, road networks, and the reports of online users. Depending on the detected refueling behavior of taxis and the estimated results of gas stations, we analyzed urban refueling behavior by comparing different groups' refueling patterns from a spatial-temporal economic view. The discovered refueling regularity can benefit a variety of applications. From the mind-set of the customer, the gas station wait time can be used to recommend the fastest choice. For government departments, they can rethink whether the current layout of stations is reasonable or whether some stations are excessively dense in an area while other areas might have a lack of infrastructure. From a business perspective, investors can analyze drivers' refueling behaviors to help choose locations that are most promising for attracting customers. We evaluated our system on a large-scale dataset, including 4 months of taxi trajectories in 2012, along with POIs and the road network in Beijing, as well as several human-collected datasets.

Further research will focus on how to determine the real-time status of gas stations. At the same time, we will make use of calculated results, the two indicators in the local and global views, to assist the design of a realistic system for application scenarios.

REFERENCES

Rekha Attri, Manvinder Pahwa, and Ashish Urkude. 2011. Brand position & customer loyalty for public sector oil marketing companies. *International Journal of Management Prudence* 2, 2, 25–35.

Linas Baltrunas, Bernd Ludwig, and Francesco Ricci. 2011. Matrix factorization techniques for context aware recommendation. In *Proceedings of the 5th ACM Conference on Recommender Systems*. ACM, New York, NY, 301–304.

Tat Y. Chan, V. "Paddy" Padmanabhan, and Seethu Seetharaman. 2007. An econometric model of location and pricing in the gasoline market. *Journal of Marketing Research* 44, 4, 622–635.

Scott M. Davis. 2000. *Brand Asset Management*. Jossey-Bass, San Francisco, CA.

Lieven De Lathauwer, Bart De Moor, and Joos Vandewalle. 2000. A multilinear singular value decomposition. *SIAM Journal on Matrix Analysis and Applications* 21, 4, 1253–1278.

Jerome H. Friedman. 2002. Stochastic gradient boosting. *Computational Statistics & Data Analysis* 38, 4, 367–378.

Yong Ge, Chuanren Liu, Hui Xiong, and Jian Chen. 2011a. A taxi business intelligence system. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, New York, NY, 735–738.

Yong Ge, Hui Xiong, Chuanren Liu, and Zhi-Hua Zhou. 2011b. A taxi driving fraud detection system. In *Proceedings of the IEEE 11th International Conference on Data Mining (ICDM'11)*. IEEE, Los Alamitos, CA, 181–190.

Chandler E. Hatton, Satish K. Beella, J. C. "Han" Brezet, and Ype Wijnia. 2009. Charging stations for urban settings: The design of a product platform for electric vehicle infrastructure in Dutch cities. *World Electric Vehicle Journal* 3, 1–13.

David L. Huff. 1964. Defining and estimating a trading area. *Journal of Marketing* 28, 3, 34–38.

Ganesh Iyer and Seethu Seetharaman. 2005. Quality and location in retail gasoline markets. In *Proceedings of the Conference on Strategic and Tactical Decision Making in Supermarket Retailing*.

Ganesh Iyer and Seethu Seetharaman. 2008. Too close to be similar: Product and price competition in retail gasoline markets. *Quantitative Marketing and Economics* 6, 3, 205–234.

Pablo Jensen. 2006. Network-based predictions of retail store commercial categories and optimal locations. *Physical Review E* 74, 3, 35101.

Alexandros Karatzoglou, Xavier Amatriain, Linas Baltrunas, and Nuria Oliver. 2010. Multiverse recommendation: N-dimensional tensor factorization for context-aware collaborative filtering. In *Proceedings of the 4th ACM Conference on Recommender Systems*. ACM, New York, NY, 79–86.

Scott Kelley and Michael Kuby. 2013. On the way or around the corner? Observed refueling choices of alternative-fuel drivers in southern California. In *Proceedings of the Transportation Research Board 92nd Annual Meeting*.

Tim Kindberg, Matthew Chalmers, and Eric Paulos. 2007. Guest editors' introduction: Urban computing. *IEEE Pervasive Computing* 6, 3, 18–20.

Ryuichi Kitamura and Daniel Sperling. 1987. Refueling behavior of automobile drivers. *Transportation Research Part A: General* 21, 3, 235–245.

Leonard Kleinrock and John Wiley. 1977. Queueing systems. *IEEE Transactions on Communications* 1, 178–179.

Kun Li, Man Lu, Fenglong Lu, Qin Lv, Li Shang, and Dragan Maksimovic. 2012. Personalized driving behavior monitoring and analysis for emerging hybrid vehicles. In *Proceedings of the 10th International Conference on Pervasive Computing (Pervasive'12)*. 1–19.

He Ling-Yun and Li Yan. 2009. Characteristics of China's coal, oil and electricity price and its regulation effect on entity economy. *Procedia Earth and Planetary Science* 1, 1, 1627–1634.

Qi Liu, Yong Ge, Zhongmou Li, Enhong Chen, and Hui Xiong. 2011. Personalized travel package recommendation. In *Proceedings of the IEEE 11th International Conference on Data Mining (ICDM'11)*. IEEE, Los Alamitos, CA, 407–416.

Craig McPherson, John Richardson, Oscar McLennan, and Geoff Zippel. 2011. Planning an electric vehicle battery-switch network for Australia. In *Proceedings of the 2011 Australasian Transport Research Forum*. 12.

Frederick W. Memmott. 1963. Home interview survey and data collection procedures. *Highway Research Record* 41, 7–12.

Olivier Monod, Albert Gaspoz, and François Golay. 2011. *Geographic Concentration of Economic Activities: On the Validation of a Distance-Based Mathematical Index to Identify Optimal Locations*. Master's Thesis. Laboratory of Geographic Information Systems, Ecole Polytechnique Federale de Lausanne, Lusanne, Switzerland.

Francesco Ricci and Bracha Shapira. 2011. *Recommender Systems Handbook*. Springer.

Chen Shaofeng. 2006. State-regulated marketization: China's oil pricing regime. *Perspectives* 7, 3, 151.

Xuan Song, Quanshi Zhang, Yoshihide Sekimoto, Teerayut Horanont, Satoshi Ueyama, and Ryosuke Shibasaki. 2013. Modeling and probabilistic reasoning of population evacuation during large-scale disaster. In *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, New York, NY, 1231–1239.

Magne Supphellen and Terri L. Rittenburg. 2001. Consumer ethnocentrism when foreign products are better. *Psychology & Marketing* 18, 9, 907–927.

Dingqi Yang, Daqing Zhang, Zhiyong Yu, and Zhiwen Yu. 2013. Fine-grained preference-aware location search leveraging crowdsourced digital footprints from LBSNs. In *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, New York, NY, 479–488.

Jing Yuan, Yu Zheng, Chengyang Zhang, Wenlei Xie, Xing Xie, Guangzhong Sun, and Yan Huang. 2010. T-drive: Driving directions based on taxi trajectories. In *Proceedings of the 18th SIGSPATIAL International Conference on Advances in Geographic Information Systems*. ACM, New York, NY, 99–108.

Nicholas Jing Yuan, Yu Zheng, Liuhang Zhang, and Xing Xie. 2012. T-Finder: A recommender system for finding passengers and vacant taxis. *IEEE Transactions on Knowledge and Data Engineering* 25, 10, 2390–2403.

Tong Zhang. 2004. Solving large scale linear prediction problems using stochastic gradient descent algorithms. In *Proceedings of the 21st International Conference on Machine Learning*. ACM, New York, NY, 116.

Yu Zheng, Yanchi Liu, Jing Yuan, and Xing Xie. 2011. Urban computing with taxicabs. In *Proceedings of the 13th International Conference on Ubiquitous Computing*. ACM, New York, NY, 89–98.

Received December 2013; revised May 2014; accepted July 2014