This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

Gaussian Mixture Models for Parking Demand Data

Tanner Fiez, *Student Member, IEEE*, and Lillian J. Ratliff, *Member, IEEE*

***Abstract*— To mitigate congestion caused by drivers cruising in search of parking, performance-based pricing schemes have received a significant amount of attention. However, several recent studies suggest location, time-of-day, and awareness of policies are the primary factors that drive parking decisions. Harnessing data provided by the Seattle Department of Transportation and considering the aforementioned decision-making factors, we analyze the spatial and temporal properties of curbside parking demand and propose methods that can improve traditional policies with straightforward modifications by advancing the understanding of where and when to administer pricing policies. Specifically, we develop a Gaussian mixture model based technique to identify zones with similar parking demand as quantified by spatial autocorrelation. In support of this technique, we introduce a metric based on the repeatability of our Gaussian mixture model to investigate temporal consistency.**

***Index Terms*— Smart parking, clustering methods, geospatial analysis, data mining.**

I. Introduction

DRIVERS cruising in search of parking is often touted to be a major contributor to congested traffic [1]. Studies have shown that the costs associated with parking-related congestion in terms of lost time, excess use of fuel, and increased pollution can be significant [2]. To combat such negative impacts, cities and researchers are examining strategies to improve parking resource management with the goal of redistributing demand more uniformly in space. These efforts can predominantly be broken into work on resource allocation and sensor management [3]–[5], object classification methods [6], [7], performance based pricing strategies [8]–[10], and behavioral models leveraging data [11]–[13]. For a comprehensive overview of smart parking solutions see [14].

The work in this paper is most closely related to that on performance-based pricing and data-driven behavioral models. However, in contrast to prior work investigating direct pricing mechanisms, we study how mechanisms such as adjusting the locations and time periods in which pricing schemes are administered can improve policy.

To motivate this line of research, we remark that there is mounting evidence supporting the basis that many factors beyond price drive parking decisions. Survey results from Los Angeles (LA) indicated that proximity to the intended destination is a more important consideration than both parking cost and time spent searching for parking [15]. On average, the respondents said the maximum distance they would be willing to park from their intended destination was just 3.07 blocks. In a study conducted in Beijing, nearly 90% of those surveyed said they chose where to park based on proximity to their destination, with less than 1% saying low price was the reason [16]. In fact, nearly 70% of respondents parked within a five minute walk to their destination.

Studies on price elasticity reinforce these findings. For example, an evaluation of the SFpark study [11] found elasticity varied immensely with location, time-of-day, day-of-week, and date of price change. Strikingly, in some cases the price elasticity was positive, which provides further evidence that price is often not the most important decision-making factor.

The inadequate awareness of parking policy is also problematic. As an example, driver behavior did not change in the SFpark study until there was an increase in marketing and advertisement during the second price adjustment. Similar outcomes are also observed in the aforementioned survey in LA, which confirmed that driver awareness of parking-related policy and technology is unquestionably low. Of those surveyed, only 31%, 24%, and 25% were aware of price changes, time-of-day pricing, and mobile applications, respectively [15].

Despite new data sources that could potentially support sophisticated management strategies, cities generally employ simple policies that benefit from being easy to track and understand. Existing policies customarily utilize static pricing schemes with morning and evening rate periods that are applied uniformly within zones containing an extensive number of block-faces. Importantly, block-faces are often grouped over long-existing regions selected prior to new technology integration or simply following existing neighborhood boundaries.

It is informative to consider properties that must hold in order for this policy form to be reasonable: (i) demand for parking needs to be similar within a zone and (ii) the spatial distribution of parking demand needs to be invariant within rate periods, over each day of the week, and from week to week. The reasons that these properties need to hold are quite clear. Applying uniform pricing in a zone with unbalanced demand induces congestion in sought-after locations, while at the same time provides no incentive for drivers to explore nearby blocks with open spots. Moreover, employing a fixed policy within rate periods at each day-of-week relies on the presumption that behavior is stationary within rate periods, independent of the day-of-week, and consistent through time.

In this work, we demonstrate that traditional policies could be improved with straightforward modifications designed to optimize for the preceding criteria. Leveraging available data sources to gain insights into the spatial and temporal properties

Manuscript received December 2, 2017; revised May 21, 2018, October 14, 2018, and April 24, 2019; accepted June 21, 2019. This work was supported in part by the NSF Grant CSN-1646912 and Grant CNS-1634136. The work of T. Fiez was supported in part by the National Defense Science and Engineering Graduate (NDSEG) Fellowship. The Associate Editor for this article was P. Kachroo. *(Corresponding author: Tanner Fiez.)*

The authors are with the Department of Electrical and Computer Engineering, University of Washington, Seattle, WA 98195 USA (e-mail: fiezt@uw.edu; ratliffl@uw.edu).

Digital Object Identifier 10.1109/TITS.2019.2939499

of parking demand estimated from paid parking transactions, we develop an approach to identify zones and time periods with similar spatial and temporal demand. Such analysis allows for simple, static pricing schemes to be more effective by providing decision makers data-informed suggestions of where and when to administer pricing schemes in terms of zones encompassing groups of block-faces and rate periods.

Specifically, we show that a Gaussian mixture model (GMM) can be used to identify groups of spatially close block-faces which have a high degree of spatial autocorrelation in observed occupancy patterns. The Bayesian information criterion (BIC) is used to select the number of zones. We supplement the model by providing a method based on the repeatability of the GMM to metricize the temporal consistency of spatial demand which governs the model, and we demonstrate through experiments that spatial demand and our model are indeed temporally consistent. We further show that the GMM zoning is consistent even with price changes. Likewise, while occupancy fluctuations can be significant from season to season, the GMM zoning remains consistent. Both of these results reaffirm that location is a primary driver of parking choice. Finally, we remark that the GMM zoning agrees with location features; e.g., blocks near tourist attractions are clustered.

This paper builds on our prior work [17] by significantly increasing the depth of analysis and experimental results. This includes exploring the effects of price changes and seasonality on both neighborhood-wide occupancy and spatial demand in support of assessing the GMM zoning approach, comparisons between GMM design choices and with auxiliary clustering methods, and the inclusion of supplementary methods for designing the spatial weight matrix when evaluating spatial autocorrelation.

The paper is organized as follows. In Section II, we describe our data sources and method for estimating demand. In Section III, we explore the spatial and temporal properties of demand as well as the effects of seasonality and price adjustments. We describe our approach using a GMM to identify zones with similar spatial demand and our method to quantify this using spatial autocorrelation in Sections IV and V, respectively. In Section VI, we present the results of our analysis using parking data from the city of Seattle and we conclude in Section VII.

II. Data Sources and Demand Estimation

We use parking transaction data, block-face supply data, and GPS location data of the block-faces from June, 2016–August, 2017 obtained via the Seattle Department of Transportation (SDOT) API.[1] We analyze a full year of data so we can examine spatial demand changes through time with respect to seasonality effects and price changes. During this time period there is nearly 14 million paid parking transactions. The paid parking transaction data includes both pay-station and mobile app-based payment records for each block-face. Paid parking is available 8AM–8PM Monday–Saturday. The supply data

[1]These data sources are all publicly available at https://data.seattle.govdata.seattle.gov.

Fig. 1. Paid parking in Belltown is broken up into the North Zone (red) and the South Zone (blue). In 2017, the North Zone was $1.00/hr between 8AM–11AM and $1.50/hr between 11AM–8PM with four hour time limits and the South Zone was $2.50/hr between 8AM–5PM and 5PM–8PM with two and three hour time limits, respectively.

consists of the estimated number of parking spaces for each block-face[2]. The GPS location data includes the latitude and longitude of both ends of a block-face.

As a proxy for demand, we use estimated parking occupancy, which is considered to be the number of active transactions divided by the supply for each block-face at each minute. Formally, the estimated occupancy of block-face i at time k is given by

$$\text{Occupancy}_i[k] = \frac{\#\text{ Active Transactions}_i[k]}{\text{Supply}_i[k]}. \tag{1}$$

While ordinarily static, variations in supply can be caused by construction, infrastructure changes, and peak hour periodic closures. Since parking prices generally do not change at any higher frequency than one hour, we aggregate the occupancies up to an hour granularity. The estimated occupancy deviates from the true occupancy because select vehicles can park for free (e.g., disabled permits), vehicles leave before the end of their paid time, and the estimated supply is inexact[2,3]. Since our analysis is based on the relative relationship between occupancies, the deviation has a negligible effect on our work.

III. Empirical Investigation of Demand Properties

In this section, we analyze the spatial and temporal properties of parking demand as well along with the effects of seasonality and price changes. We focus on the Belltown neighborhood in Seattle, WA which is a rapidly growing mixed-use area with both the highest population density [18] and the most complete coverage of curbside parking across Seattle neighborhoods. In total, the neighborhood contains over 250 paid block-faces. These characteristics make Belltown an ideal for a case study. Belltown parking policy is described in the caption of Fig. 1.

Since the data needed to facilitate empirical investigations of demand trends resulting from the behavior of drivers looking to park has only recently become available, there has been limited study of this matter. To that end, this section contributes data-informed observations regarding demand trends. Notably, we find evidence supporting the basis that the primary decision factors influencing parking decisions are location, time-of-day, and day-of-week. This conclusion provides motivation for our

[2]Parking spaces in Seattle are not marked; the number of spaces in a block-face is estimated by dividing the legal parking zone into 25ft increments.

[3]These factors can cause the estimated occupancy to exceed 100%; we clip the maximum estimated occupancy at 150% which occurs at less than 0.45% of the hourly occupancy instances over all block-faces.

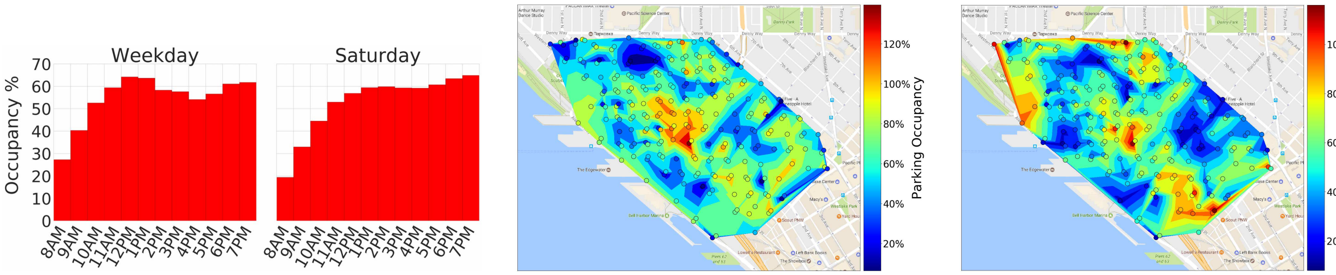

(a) (b) Friday 7PM (c) Saturday 11AM

Fig. 2. (a) Mean occupancies over Summer 2017 (June–August) in Belltown for a weekday (Wednesday) and the weekend day (Saturday). (b–c) Contours of the mean occupancies in Belltown within Summer 2017 at Friday 7PM and Saturday 11AM, respectively. Each scatter point is the midpoint of a block-face.

focus on how these attributes can be further scrutinized to improve simple policy with direct adaptations.

A. Temporal Properties

In Fig. 2a, the mean occupancy profiles over the entire Belltown neighborhood are shown at each hour paid parking is available for both Wednesday, a representative weekday, and Saturday. Occupancy profiles for Monday–Friday are similar, while the occupancy profile for Saturday follows a different trend. During weekdays, occupancy increases from the start of paid parking until demand peaks near lunch time. It then decreases during the afternoon, until there is another peak in the evening when people tend to have dinner. In contrast to weekdays, on Saturday the demand for parking nearly continuously increases throughout the day. These observations highlight that a reasonable parking policy may use unique weekday and weekend pricing schemes and consider the temporal properties of demand that can be driven by neighborhood features such as the presence of certain business types.

B. Spatial Properties

In Fig. 2b, the occupancy at 7PM on Friday is nearly uniformly distributed throughout Belltown, with the exception that there is an area of much higher occupancy in the center of the neighborhood that we conjecture is driven by the high density of bars and restaurants present there. Interestingly, this area also appears to be up against the divide of the north and south paid parking zones—denoted by the red and blue block-faces respectively in Fig. 1—which have a \$1.00/hr price difference at this time. One could posit that an improved division of paid parking zones, such that this area was solely encompassed by a zone, could reduce the congestion there.

In Fig. 2c, the occupancy at 11AM on Saturday has a more diverse distribution, but most importantly the areas of high occupancy, with the exception of a few block-faces in the center of the neighborhood, are in very different locations. The source of the high occupancy areas is immediately clear, as the top and bottom of the neighborhood are the closest parking to famous weekend tourist attractions in Seattle. Just above the top of the neighborhood is the Space Needle, and just below the bottom of the neighborhood is Pike Place Market.

The important takeaway from this example is that depending on the time-of-day and day-of-week, certain areas are more desirable than others; consequently, demand is not uniform within the neighborhood. Yet, by and large, these areas are being priced identically to locations that are not nearly as coveted. Hence, the properties we outlined in Section I that must hold for a static policy to be effective are being violated. Namely, uniform pricing is being applied to zones with dissimilar spatial demand, and invariant policy is being employed within rate periods and at each day of the week despite the fact that the behavior is neither stationary within the rate periods nor is it independent of the day-of-week.

C. Effects of Seasonality

Parking demand does indeed exhibit fluctuation between seasons as indicated in Fig. 3. We find that on average the occupancy of block-faces increases or decreases by approximately 10% between all seasons. Yet, the percentage of block-faces that increase between seasons is heavily dependent on the seasons being transitioned between. From Summer to Fall 2016, and similarly between Fall 2016 to Winter 2017, more block-faces decrease in occupancy than do increase at the majority of time periods. Intuitively, this makes sense. During the day people follow their regular routines—e.g., parking for work—while in the evening, shorter days and worse weather have the effect of causing people to become less likely to go out to businesses and participate in activities.

There is an analogous trend between Winter to Spring 2017 and Spring to Summer 2017. Between these seasons, at most of the paid parking times, more block-faces increase in occupancy than do decrease. During the morning hours when people follow their normal routines, the percentage of block-faces whose occupancy increases is often at or just above 50%. At times in which demand may be driven more by businesses, such as near lunch and in the evening, a higher percentage of block-faces increase in occupancy. The intriguing property of these observations, which we discuss further in Section VI-E, is that despite variations between seasons, the way occupancy is distributed (spatial demand) does not vary significantly. This suggests that static policy schemes considering location can be robust to the effects of seasonality.

D. Effects of Price Changes

In Seattle, parking prices change each year after an annual parking study is conducted. In July, 2016 the price to park in the North Zone of Belltown (red block-faces in Fig. 1) decreased from \$1.50/hr to \$1.00/hr in 8AM–11AM. To investigate the impact the price change had on parking behavior, we examine the month before the price change, June, 2016, and the month one year following the price change, June, 2017.

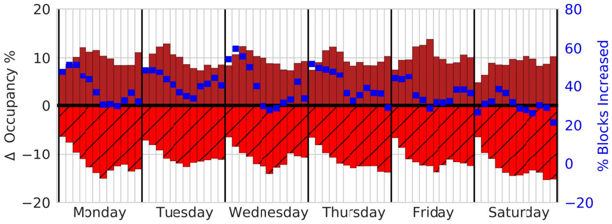

(a) Summer (Jun.–Aug.) to Fall (Sept.–Nov.) 2016.

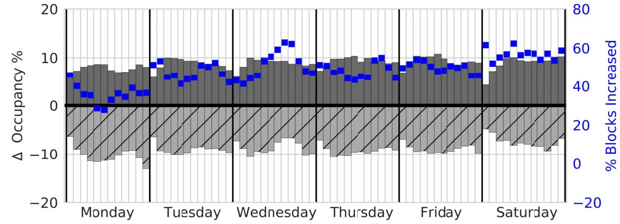

(b) Fall (Sept.–Nov.) 2016 to Winter (Dec. '16–Feb. '17) 2017.

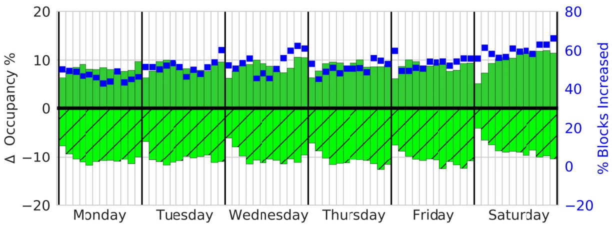

(c) Winter (Dec. '16–Feb. '17) 2017 to Spring (Mar.–May) 2017.

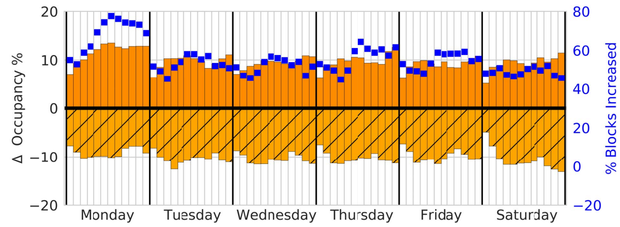

(d) Spring (Mar.–May) 2017 to Summer (Jun.–Aug.) 2017.

Fig. 3. Seasonality effects on demand in Belltown: each bar indicates an hour of the active paid parking times in a day between 8AM—8PM. The bars above the x-axis indicate the mean percentage increase in occupancy at each hour of available paid parking for the block-faces that saw increased occupancy between seasons. Similarly, the hatched bars below the x-axis indicate the mean percentage decrease. The mean variation between seasons in either direction is 10%. The blue squares indicate the percentage of block-faces that increased in occupancy between seasons.

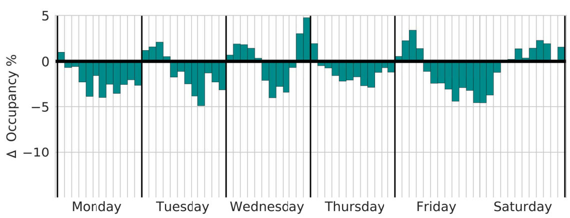

(a) North Zone in Belltown.

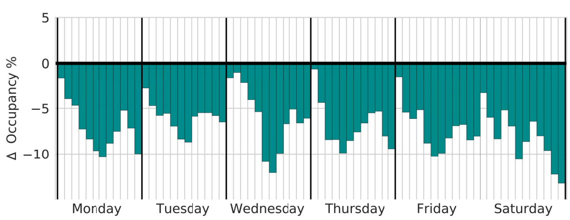

(b) South Zone in Belltown.

Fig. 4. Occupancy change in North Zone (Fig. 4a) and South Zone (Fig. 4b) from before (June, 2016) to after (June, 2017) the Belltown price change.

Fig. 4 shows the relative change in occupancy from before the price change to a year after the price change in the North Zone (Fig. 4a) and the South Zone (Fig. 4b). Interestingly, the times that see an increase in occupancy in the North Zone primarily occur within the time interval in which the price was decreased. However, the shift is rather insignificant. Occupancy went down consistently in the South Zone[4]. It is worth noting though that the least significant declines in occupancy occur in the time interval of the price change in the North Zone. This indicates that the price change is likely not the primary factor causing the demand to rise in the North Zone. It also confirms prior work showing elasticity to price is mixed and further reaffirms that pricing alone may be limited in its efficacy to change driver behavior.

The mean variation in occupancy over the entire year for block-faces is approximately 20%, whereas between seasons we observed approximately a 10% mean variation (see Fig. 3). Commensurate with variation between seasons, we find (and show in Section VI-E) that spatial demand does not vary significantly before and after the price change.

IV. Gaussian Mixture Model

The preliminary data analysis of the previous section reinforces that demand for parking is guided by location, time-of-day, and day-of-week. With this in mind, we focus on the question of how to zone block-faces at each hour-of-day and day-of-week to reflect the spatial demand. Specifically, we model spatial demand with a GMM by using it as an unsupervised clustering method to find zones containing block-faces that are spatially close and have similar demand at a given time. Then, by observing the variation in the GMM across each time point, we are able to determine a relationship between the zones discovered by the GMM at each hour-of-day and day-of-week as described in Section VI-D.

Alternative methods such as k-means clustering, spectral clustering, and hierarchical clustering [19] may be considered; however, with respect to the problem at hand, we believe a GMM is a natural choice since it is a probabilistic method that has simple design choices with principled approaches.

We do not desire an exact recovery of an underlying clustering since one does not exist; rather, we seek a clustering that best explains the data with the understanding that the data points corresponding to block-faces are not clearly separable. This is to say that we would like to capture some of the

[4] Historically in Seattle utilization of on-street parking has consistently risen. The decline in demand is a new trend (per personal communication with SDOT). A possible cause of this may be the recently expanded light rail system and notable increase in use of ride- and bike-sharing services.

uncertainty inherent to the problem. Probabilistic clustering methods naturally lend themselves to such a task as they make soft-assignments that provide a notion of confidence that a data point belongs to a cluster and they incorporate information from all data points in forming clusters rather than from only data points contained within a cluster.

GMMs also support application-informed design choices. Paramount of which to our task is that the structure of the covariance matrix can be selected to reflect the geometry of the underlying problem. Moreover, there are principled approaches to select the number of clusters (see Section IV-B). On the other hand, spectral and hierarchical clustering both require several design choices that can significantly impact the learned model and accepted methods for selecting the number of clusters rely on arguably inferior heuristics.

We provide a method in Section IV-C to assess the feasibility of GMM zone-based policy that evaluates whether demand is comparable from week to week and if a model learned from historical data is reasonable for the future bearing in mind demand variations. This method relies on the GMM being an inductive clustering algorithm, meaning a function on the input space is induced that can map any data point to a label.

A. Model Description

The GMM is a probabilistic method to model a distribution of data with a mixture of multivariate Gaussian distributions, each with weight π_j, mean vector μ_j, and covariance matrix Σ_j. The probability distribution of each sample data point x_i in a GMM with k mixture components is given by

$$p(x_i|\pi, \mu, \Sigma) = \textstyle\sum_{j=1}^{k} \pi_j \mathcal{N}(x_i|\mu_j, \Sigma_j). \quad (2)$$

Each sample $x_i = \begin{bmatrix} x_{i,\text{latitude}} \; x_{i,\text{longitude}} \; x_{i,\text{occupancy}} \end{bmatrix} \in \mathbb{R}^3$ is a vector containing both location and demand features for a block-face at a certain time, or an aggregated time period for demand (e.g., mean demand at an hour-of-day and day-of-week over a period of weeks). The data provided to the model is a matrix of n samples, corresponding to the number of block-faces input, stacked as $x = \begin{bmatrix} x_1 \cdots x_n \end{bmatrix}^\top \in \mathbb{R}^{n\times 3}$. In our implementation we normalize features column-wise to be in $[0, 1]$. A motivation for the features we choose is their simplicity and the availability of data; they exactly capture the spatial demand aspects of the data we have and work to trade-off grouping block-faces that are close and that have similar demand. It is also possible that these features capture unobserved information or effects of latent variables such as the price, type of area (residential, commercial), type of parking (parallel, angled), etc.

Let $z = \begin{bmatrix} z_1 \cdots z_n \end{bmatrix}$ be a n–dimensional vector of indicator variables which are the latent component labels for samples. The prior on the probability of a sample belonging to a mixture component can then be expressed as $p(z_i = j) = \pi_j$. The parameter π must satisfy the restrictions that $\pi_j \in [0, 1]$ and $\sum_{j=1}^{k} \pi_j = 1$. The likelihood of a sample belonging to a mixture component is given by $p(x_i|z_i = j) = \mathcal{N}(x_i|\mu_j, \Sigma_j)$, where the multivariate Gaussian distribution is

$$\mathcal{N}(x_i|\mu_j, \Sigma_j) = \frac{\exp\left(-\frac{1}{2}(x_i-\mu_j)^\top \Sigma_j^{-1}(x_i-\mu_j)\right)}{(2\pi)^{\frac{3}{2}}|\Sigma_j|^{\frac{1}{2}}}. \quad (3)$$

The form of the covariance matrix Σ determines the shape of the Gaussian level set ellipses, and is the primary design choice. It may be desirable to enforce that the zones proposed from a model reflect the city grid when it has a regular structure since such a constraint can lead to an easier to understand policy. To learn a model guaranteed to reflect the geometry of the underlying problem and follow a parametric structure, we can consider a diagonal covariance matrix. The diagonal covariance structure constrains the ellipses to be independent in each dimension and aligned with the axes of the data. To align the spatial axes of our data with a city grid, we can apply a simple coordinate transformation using a rotation matrix since our coordinate axes are the cardinal directions. This means we obtain updated spatial coordinates as follows:

$$\begin{bmatrix} x'_{i,\text{latitude}} \\ x'_{i,\text{longitude}} \end{bmatrix} = \begin{bmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{bmatrix} \begin{bmatrix} x_{i,\text{latitude}} \\ x_{i,\text{longitude}} \end{bmatrix}, \quad (4)$$

where θ is the angle counter-clockwise of north that a grid is aligned. We transform the spatial coordinates this way with an estimated angle of $\theta = 35$ when using the diagonal covariance structure in our experiments for Belltown.

Furthermore, we consider a full covariance structure, which allows the ellipses to be arbitrarily oriented. This choice is of interest since it can provide the best representation of the data as it is unconstrained. It can also be a reasonable selection when a city grid is highly non-regular or if there is no requirement on the zones reflecting the city grid precisely. The spatial coordinates do not need to be transformed for this method. We explore the impact of the covariance structure on the GMM and our evaluation metrics in Section VI.

The objective function of the GMM is the log likelihood of the data given by

$$\text{LL} \triangleq \log \; p(x|\pi, \mu, \Sigma) = \textstyle\sum_{i=1}^{n} \log \sum_{j=1}^{k} \pi_j \mathcal{N}(x_i|\mu_j, \Sigma_j).$$

We employ the expectation-maximization (EM) algorithm [20] to optimize this objective. The EM algorithm consists of an initialization of the unknown parameters and two steps—the expectation step (E-step) and the maximization step (M-step)—which are repeated until convergence as determined by checking at each iteration if the change in the log likelihood (ensured to be positive [21]) between iterations is smaller than some sufficiently small ϵ: $\Delta\text{LL} \triangleq \text{LL}^i - \text{LL}^{i-1} \leq \epsilon$.

In the E-step, given the current parameter values, the expected values of the unobserved component labels are computed. These are the posterior probabilities which capture the responsibility that component j takes for data point i [19]. Formally, we denote this term as

$$r_{i,j} \triangleq p(z_i = j|x_i, \pi_j, \mu_j, \Sigma_j) = \frac{\pi_j \mathcal{N}(x_i|\mu_j, \Sigma_j)}{\sum_{j'=1}^{k} \pi_{j'} \mathcal{N}(x_i|\mu_{j'}, \Sigma_{j'})}. \quad (5)$$

In the M-step, the parameter values are updated to maximize the log likelihood. Hence, for $1 \leq j \leq k$, we update

π_j, μ_j, Σ_j as follows:

$$\left.\begin{aligned} \pi_j &= \frac{1}{n}\boldsymbol{r} \\ \mu_j &= \frac{1}{\boldsymbol{r}}\sum_{i=1}^{n} x_i r_{i,j} \\ \Sigma_j &= \tfrac{1}{\boldsymbol{r}}\Big(\sum_{i=1}^{n}(x_i-\mu_j)^\top(x_i-\mu_j)r_{i,j}\Big) \end{aligned}\right\} \tag{6}$$

where $\boldsymbol{r} = \sum_{i=1}^{n} r_{i,j}$.

Once the convergence criteria is met, we make assignments of each sample x_i to the component label j which maximizes the responsibility $r_{i,j}$, meaning $z_i^* = \arg\max_j \ r_{i,j}$.

The objective function is non-convex, which only guarantees that we find local minima. Hence, we run the algorithm for several random initializations and retain the model from the iteration that resulted in the highest log likelihood.

B. Model Selection

The model selection problem for a GMM is to determine the number of mixture components in the model. We use the Bayesian information criterion (BIC) score [22] for this purpose. It is given by

$$\mathrm{BIC} = -2\mathrm{LL} + \log(n)\nu, \tag{7}$$

where n is the number of data points and ν is the number of degrees of freedom in the GMM. For a GMM containing k components with d dimensional data, $\nu = k(2d+1)$ and $\nu = k(d+d^2+1)$ for diagonal and full covariances, respectively.

To determine the number of components in our experiments, we performed a search over k and evaluated the mean BIC of the GMMs learned on each day of the week and hour of the day combination paid parking is available using the mean occupancies at each instance. We then selected the value of k that minimized the mean BIC.

C. Consistency Metric

Given that in Sections III-C and III-D we observed there is a non-trivial variation in the demand at block-faces through time, we want to find how consistently spatial demand is distributed. As alluded to in our introduction, stationary policy relies on the premise that spatial demand is consistent from week to week at a given time-of-day and day-of-week. Indeed, weekly variations in spatial demand should not radically alter what the best zoning of block-faces is according to learning method used to determine such a zoning. This can hold if the variations in spatial demand are reasonable and when the learning method is not overly sensitive to perturbations. In consideration of this, we propose a consistency metric to assess the feasibility of GMM zone-based policy that evaluates whether demand is comparable through time and if a model learned from past data can be reasonable for the future. The procedure to determine the consistency metric value under the GMM at a day of the week and hour of the day is as follows:

Step 1 : For the chosen day of the week and hour of the day, select a specific date and learn a GMM using the occupancy data at this instance.

Step 2 : Assign component labels to each block-face for all other instances with the same day of the week and hour of the day using the learned model.

Step 3 : Determine the percentage of block-faces which were assigned to the same component as they were in the original GMM that was learned.

Step 4 : Repeat **Step 1**–**Step 3** switching the date on which the GMM is learned, and then average over the percentages computed at each iteration.

We explore this method and discuss the results in Section VI.

V. Spatial Autocorrelation

For the simple policy form we have been discussing to be effective, it is imperative that parking demand is similar and correlated between block-faces within zones. The reason is that if demand bears no spatial similarity or relation, then it would be misguided to carry out policy uniformly through space.

We use a standard measure of spatial autocorrelation known as Moran's I [23] to quantify the degree of spatial similarity or relation present in the demand data. Moran's I is defined by

$$I = \frac{n}{\sum_i\sum_j w_{i,j}}\frac{\sum_i\sum_j w_{ij}(o_i-\bar{o})(o_j-\bar{o})}{\sum_i(o_i-\bar{o})^2}, \tag{8}$$

where n denotes the number of block-faces, o_i denotes the occupancy for block-face i, $\bar{o}$ denotes the mean occupancy over all block-faces, and $W = (w_{i,j})_{i,j=1}^{n}$ is a matrix of spatial weights with zeros along the diagonal.

Values of I range from -1 (indicating perfect dispersion) to 1 (indicating perfect clustering of similar values). The Moran's I value can be used to find a $z-$score and then a $p-$value to determine whether the null hypothesis, that the data is randomly disbursed, can be rejected.

We are interested in assessing the spatial autocorrelation locally and globally in a neighborhood or region, within currently designated paid parking areas, and within the zones we find with the GMM. The spatial weight matrix W can be designed to evaluate each of these objectives as we describe in the following subsections. Several of these methods require a distance metric to determine the spatial relationship between block-faces and we employ the Manhattan (ℓ_1) distance metric. Importantly, we transform our spatial coordinate system to align with the city grid as in (4) so that the distance reflects a path that can be followed along the grid.

In Section VI, we evaluate each method of creating the spatial weight matrix by determining whether the p-values are significant using a two-sided p-value with a significance measure of .01. We report the percentage of instances—given by the occupancy at a date and time—that are significant.

A. Assessing Local and Global Spatial Autocorrelation

As indicated in Section III, parking demand displays spatial diversity when considering broad neighborhoods. A logical follow-up question is whether there is (at least) local similarity or relation. To evaluate this objective, we create the weight matrix by setting values of w_{ij} to one if block-face j is of the k nearest neighbors to block-face i in terms of the Manhattan distance in the rotated GPS coordinate space and zero if it is not. We experiment using a range of values for k in order to evaluate what size of local neighborhood contains significant spatial autocorrelation. In order to take a more global view of the spatial demand we also experiment creating the weight matrix by using a distance based metric. That is, we set values of w_{ij} to be the Manhattan distance between block-face i and j in terms of the rotated GPS coordinates, normalized in [0, 1] with weight one given to the closest block-face to i and weight zero given to the furthest block-face from i.

B. Assessing Spatial Autocorrelation in Current Zones

We are also interested in the spatial autocorrelation within the currently designated paid parking zones by the city of Seattle. This will help us assess current policies and provide a means to make comparisons with our method of selecting paid parking zones. To measure the spatial autocorrelation within the current zones we create the weight matrix by setting values of w_{ij} to one if block-faces i and j are in the same parking zone and zero if they are not.

We also investigate a distance based metric within the paid parking zones in a similar manner as described in Section V-A. In this method we set values of w_{ij} to be the Manhattan distance between block-face i and block-face j in terms of the rotated GPS coordinates, normalized in [0, 1] with weight one given to the closest block-face to i in the same paid parking zone and weight zero given to the furthest block-face from i in the same paid parking zone. Block-faces in different paid parking zones are given a weight of zero.

C. Assessing Spatial Autocorrelation in GMM Components

One of the goals of the GMM approach is to identify groups of block-faces that are spatially close and have similar demand patterns. This can also be interpreted as finding zones of connected block-faces where there is significant spatial autocorrelation. To gauge our success in doing so, we create W by setting values of w_{ij} to one when block-faces i and j are in the same GMM component and zero otherwise.

We also use a distance based metric for this objective. We set values of w_{ij} to be the Manhattan distance between block-face i and block-face j in terms of the rotated GPS coordinates, normalized in $[0, 1]$ with weight one given to the closest block-face from i assigned to the same mixture component and weight zero given to the furthest block-face from i assigned to the same mixture component. Block-faces assigned to different mixture components are given a weight of zero.

VI. Experiments and Results

In this section, we present consistency and spatial autocorrelation results for the GMM, provide insights into the spatial and temporal properties of demand, and examine the impact of seasonality and price adjustments on the GMM.

A. Modeling Belltown With a GMM

In Fig. 5, we provide a representative comparison between the GMM learned when using a diagonal (Fig. 5a) and full (Fig. 5b) covariance matrix, respectively. The results of our evaluation criteria, in terms of the consistency metric and spatial autocorrelation, are similar for the diagonal and full covariance designs. The confidence ellipses in the case of the diagonal covariance structure are more aligned with the city grid than from the full covariance structure as we would expect since they are constrained this way. However, we find that in general the GMM with the full covariance structure does manage to produce zones that reflect the city grid to a reasonable approximation. This is to say that the GMM with the full covariance structure that has a maximum amount of freedom ends up uncovering some of the structure present in the problem, providing support for approaches that are more data-driven. We find this to be an interesting result that underlies the role the grid structure plays in shaping spatial demand. Depending on the policy maker, each choice of covariance may be considered and therefore we present examples and evaluation results for each method. In general, we find that the full covariance method produces zones that appear to be more reasonable owing to the flexibility it has compared to the diagonal covariance method.

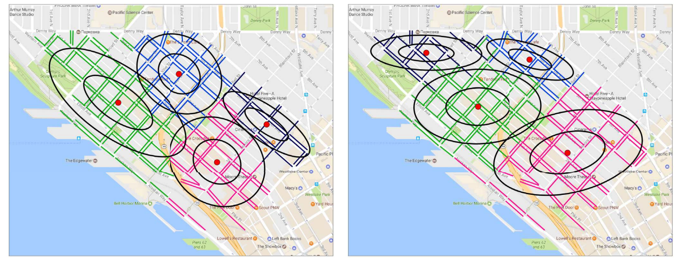

(a) Diagonal Covariance (b) Full Covariance

Fig. 5. Comparison of GMM in Belltown with diagonal and full covariance structures using mean occupancies within Summer 2017 at Wednesday 9AM.

Notably, it is clear that with each covariance structure we are able to find separable zones in which block-faces spatially close are included in the same mixture components. This is important due to the fact that while there may be spatial demand diversity in Belltown, we are able to find zones in which block-faces have similar demand thereby validating that zone-based pricing is viable. In experimental comparisons, k-means tended to exhibit the highly undesirable property of discovering zones composed of spuriously located block-faces rather than contiguous groups. Fig. 6c provides a representative example of this observation.

Recall Figs. 2b and 2c, which illustrate the mean spatial demand in Belltown within Summer 2017 at Friday 7PM and Saturday 11AM. Figs. 6a and 6b provide an example use of the GMM with full and diagonal covariance matrix structures respectively using the same data. The examples indicate that the model we learn is related to the time-of-day and day-of-week. For example, the model we learn for Friday night is quite different from the model we learn for Saturday morning. This observation asserts that the spatial component of demand depends on the temporal component. We discuss further spatial and temporal insights in Section VI-D.

Using the BIC described in Section IV-B, we find four mixture components—corresponding to four paid parking zones—to be optimal. Yet, at present, Belltown has two paid

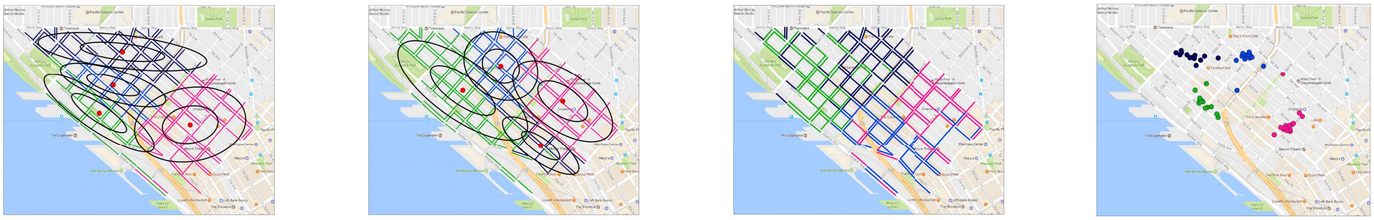

(a) Full Σ, Friday 7PM (b) Diag Σ, Saturday 11AM (c) k-means, Saturday 11AM (d) Full Σ Spatial Centers

Fig. 6. GMMs with four components learned using the mean occupancies in Belltown within Summer 2017 at Friday 7PM with a full covariance matrix in (a), and at Saturday 11AM with a diagonal covariance matrix in (b). Block-faces are colored by the mixture component label. The ellipses indicate the first and second standard deviations of the component's GPS coordinates, with the red scatter points indicating their centers. (c) Block-face labels designated via k-means for Saturday 11AM in Summer 2017; zones from this method fail to contain contiguous groups of block-faces. (d) Clustering of the spatial centers of the mixture components from the GMMs with full covariance structures learned on all Fridays at 11AM in Summer 2017.

TABLE I

MEAN CONSISTENCY METRIC RESULTS FOR BELLTOWN PAID PARKING HOURS OVER EACH PAID PARKING DAY FOR SUMMER 2017. UNITS ARE %

Method\Hour	8AM	9AM	10AM	11AM	12PM	1PM	2PM	3PM	4PM	5PM	6PM	7PM	**Mean**
GMM Diag	72.4	80.8	86.6	88.1	88.2	87.8	86.9	87.4	84.8	83.9	84.4	84.0	**84.6**
GMM Full	72.7	83.0	87.3	88.4	87.7	87.9	87.2	86.6	86.6	86.0	86.2	85.3	**85.4**
k-means	81.3	79.7	78.8	77.1	75.1	75.3	76.0	76.4	77.8	77.7	76.5	73.1	**77.1**

TABLE II

MEAN CONSISTENCY METRIC RESULTS FOR BELLTOWN PAID PARKING DAYS OVER EACH PAID PARKING HOUR FOR SUMMER 2017. UNITS ARE %

Method\Day	M	T	W	Th	F	S
GMM Diag	83.7	85.2	85.9	86.8	86.0	80.1
GMM Full	85.4	85.5	86.2	86.7	86.3	82.2
k-means	77.9	78.2	78.2	77.7	77.3	73.1

parking zones, suggesting there is potential to improve the existing policy design methods for selecting the number of zones. However, we mention that with two mixture components the GMM does learn zones similar to those designated by SDOT.

B. Consistency Metric Results

The results in Tables I and II establish that spatial demand is reasonably consistent from week-to-week at a fixed time, and that the GMM is robust to modest demand variations. With the exception of the first hour of paid parking, the consistency value for each hour of the day is high, ranging from 80.8%–88.2% for the GMM with a diagonal covariance matrix and from 83.0%–88.4% for the GMM with a full covariance matrix. Furthermore, on a given day of the week, the consistency values range from 80.1%–86.8% with a mean of 84.6% for the GMM with a diagonal covariance matrix and from 82.2%–86.7% with a mean of 85.4% for the GMM with a full covariance matrix. The mean daily consistency values at each other season we analyze were nearly the same for the GMM with a diagonal covariance matrix, which had corresponding values of 86.2%, 85.0%, 84.3%, and 84.5% for Summer 2016, Fall 2016, Winter 2017, and Spring 2017, respectively. Similarly, for the GMM with a full covariance matrix, the values were 84.5%, 83.1%, 84.0%, and 85.1%. In comparison, k-means is significantly more sensitive to spatial demand variations with nearly 10% lower consistency on average. The GMM consistency results suggest that using historical trends with our model for policy design is reasonable.

Focusing on the GMM, we investigate how the spatial centers of the mixture components change each week; e.g., given a day-of-week and hour-of-day, we use k-means clustering on the spatial centers that were found at each date with the same corresponding day-of-week and hour-of-day. Fig. 6d shows an example of clustering the centers from full covariance GMMs learned on each Friday at 11AM in Summer 2017. The centers are tightly clustered with limited change in location. By finding the centroids of each of the k-means clusters, and calculating the mean distance from each centroid to the points in that respective cluster, we can describe this change in terms of distance. In Fig. 6d, we find the mean distance of the points to their respective centroids to be just 61 meters (m). The mean of the values at each hour-of-day and day-of-week is 92m. In Summer 2016, Fall 2016, Winter 2017, Spring 2017 the results are comparable with values of 88m, 83m, 77m, and 77m, respectively. For the diagonal covariance matrix, the matching values are 77m, 79m, 73m, 74m, and 67m.

C. Spatial Autocorrelation Results

The spatial autocorrelation results provide a way of quantifying improvement of the GMM approach over current paid parking zones. Using spatial autocorrelation as a metric allows for assessment of spatial consistency and thus how reasonable it is to implement a spatially uniform policy in a particular area. Table III summarizes the results of the spatial autocorrelation analysis in Belltown for each season and method of creating the spatial weight matrix. The results are reported as follows: (rows 1–4) assessing the local and global spatial autocorrelation as described in Section V-A; (rows 5–6) assessing the spatial autocorrelation in the current paid parking zones as described in Section V-B; (rows 7–10) assessing the spatial autocorrelation in the GMM components with diagonal (rows 7, 9) and full (rows 8, 10) covariance structures as described in Section V-C.

The k-nearest neighbor method from Section V-A reveals that within Belltown there is almost always significant local

TABLE III

SPATIAL AUTOCORRELATION RESULTS FOR BELLTOWN IN EACH SEASON AND METHOD OF CREATING THE SPATIAL WEIGHT MATRIX. THE VALUES INDICATE THE PERCENTAGE OF INSTANCES IN OUR DATA SET IN WHICH THE p-VALUE OF MORAN'S I WAS SIGNIFICANT

	Method\Season	**Summer '16**	**Fall '16**	**Winter '17**	**Spring '17**	**Summer '17**	**Mean**
Sec. V-A	$k = 3$	57.1	85.4	83.0	91.0	95.1	**82.3**
	$k = 5$	75.7	92.9	91.3	95.6	98.9	**90.9**
	$k = 10$	86.6	97.1	96.8	99.1	99.8	**95.9**
	Distance	29.9	55.6	56.2	48.3	26.1	**43.2**
Sec. V-B	Area Connect	53.1	68.1	65.0	65.6	59.8	**62.3**
	Area Distance	62.3	82.8	81.1	83.4	78.4	**77.6**
Sec. V-C	GMM Diag Connect	97.5	99.7	99.7	99.7	100.0	**99.8**
	GMM Full Connect	99.0	99.9	99.9	99.6	100.0	**99.7**
	GMM Diag Distance	97.8	99.7	99.7	99.8	100.0	**99.8**
	GMM Full Distance	99.9	100.0	100.0	99.8	100.0	**99.9**

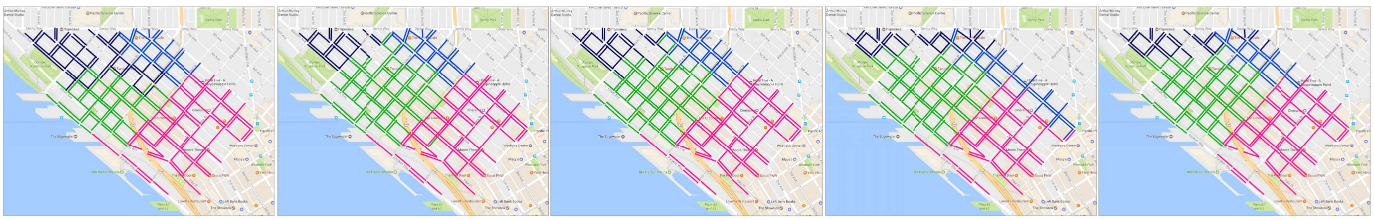

Fig. 7. GMM zoning (full covariance) for Summer 2017 weekdays at 9AM using mean occupancies. Each figure corresponds to a day Monday–Friday.

spatial autocorrelation, indicating block-faces adjacent to each other see similar demand properties. The distance based method from Section V-A confirms our previous observations that there is neighborhood-wide diversity of demand within Belltown as the frequency of significant spatial autocorrelation is much lower. The methods assessing the current paid parking zones from Section V-B demonstrate that it is often not the case that there is significant spatial autocorrelation within these zones which suggests that static pricing policies can be improved by considering location. The methods assessing the zones learned using the GMM from Section V-C demonstrate that our approach provides meaningful improvement. The results show that at nearly all times the spatial autocorrelation within the mixture components is significant. Finally, we remark that the results are similar for each covariance design choice with the GMM.

The zones we find also have lower variance in the demand than the current paid parking zones in Belltown. The mean occupancy variance in the GMM zones over each paid parking time ranges from 4.9% to 6.3% with a mean of 5.5% over the seasons with the diagonal covariance matrix, and from 5.9% to 6.7% with a mean of 6.1% over the seasons with the full covariance matrix. This is in contrast to the occupancy variance in the current paid parking zones, which ranges from 9.0% to 10.1% with a mean of 9.5% over the seasons.

D. Spatial and Temporal Insights

The GMM serves as a means to determine how to zone block-faces at a given time to reflect the spatial demand. However, the relationship between the zones discovered at each hour-of-day and day-of-week is important to guide rate periods and give insights into the how spatial demand evolves over each day and the week. As it turns out, visually examining the zones suggested from the GMM at each given time provides us with this information. The following insights are with respect to the GMM with a full covariance structure. We can draw vaguely similar conclusions for the diagonal covariance structure, but the relations are not as clear.

Notably, we find that Monday–Friday (weekdays) from 8AM–4PM similar models are learned. Likewise, we find that weekdays from 4PM–8PM similar models are learned, which are different from those learned weekdays from 8AM–4PM. We observe that models we learn for Saturday are unique and need to be considered on their own. Fig. 7 serves to show what we observe by considering a specific hour in the 8AM–4PM interval at each day of the week. To give an idea of what the model generally is like weekdays from 4PM–8PM we refer the reader to Fig. 6a. The preceding observations indicate that, based off of our model, it would make the most sense to have two weekday pricing periods—8AM–4PM and 4PM–8PM—for the zones we commonly find at these respective time periods, and a unique Saturday pricing scheme. Interestingly, these observations mirror much of what we discovered in Section III-A and III-B. The results are also compelling because they are quite different than the policies in place now, while still being surprisingly simple. Currently, the pricing periods in Belltown are 8AM–11AM and 11AM–8PM; no individual consideration is given to Saturdays.

In comparing our model to the zones in place at present, we note key similarities and differences. At nearly all paid parking times during weekdays, we learn a mixture component that covers a zone similar to that of the South Zone in Belltown. Yet in the North Zone of Belltown, our model typically learns to divide up what is now the North Zone into three distinct zones. This implies that the South Zone may be simple enough to consider as is, while improvements can be made to policies in the North Zone.

E. Seasonal and Price Changes

In previous sections, we have seen that the consistency, spatial autocorrelation, and variance results in Belltown are similar between seasons. Likewise we find that the GMM

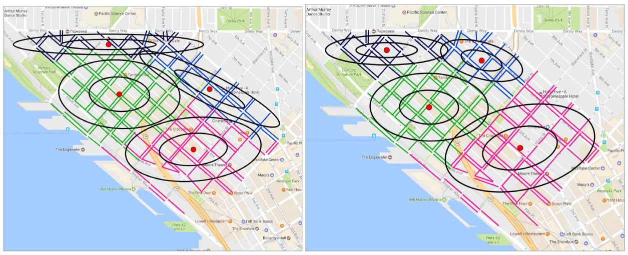

(a) Wednesday 10AM—June, 2016 (b) Wednesday 10AM—June, 2017

Fig. 8. Full covariance GMM analysis before (Fig. 8a) and after (Fig. 8b) the price change in Belltown using the mean occupancies in the respective months.

we learn within the different seasons coincide. These results indicate that in terms of the spatial demand, the variation between seasons is insignificant, despite that in Section III-C we showed the variation in occupancy is non-negligible. Thus, we conclude that demand does indeed fluctuate over time, but the fluctuation in the relative levels of demand between block-faces is inconsequential, implying static policies can be robust to seasonal variation effects. Analyzing the price change in 2016 in Belltown we draw similar conclusions. While occupancy decreased from the month before the price change to the month one year following the price change, the models we learn using data from the two time periods hardly change. Fig. 8 shows an example of this. These results corroborate studies of SFPark suggesting price control methods may not produce the desired behavior changes [11].

VII. Discussion and Future Work

We provide an in depth analysis of the spatial and temporal properties of parking demand using real data, as well as an interpretable way to find zones of connected block-faces which exhibit a high degree of spatial autocorrelation using a GMM. Furthermore, we establish that the GMM are not sensitive to modest variations in the spatial distribution of demand and as a result are consistent through time. Our analysis focused on Belltown as a case study, however we provide analysis for a broader area in [24].

References

[1] D. C. Shoup, "Cruising for parking," *Transp. Policy*, vol. 13, no. 6, pp. 479–486, 2006.

[2] D. Shoup and H. Campbell, *Gone parkin'*, vol. 29. New York, NY, USA: The New York Times, 2007.

[3] Y. Geng and C. G. Cassandras, "A new 'smart parking' system based on optimal resource allocation and reservations," *IEEE Trans. Intell. Transp. Syst.*, vol. 14, no. 3, pp. 1129–1139, Sep. 2013.

[4] A. O. Kotb, Y.-C. Shen, X. Zhu, and Y. Huang, "iParker—A new smart car-parking system based on dynamic resource allocation and pricing," *IEEE Trans. Intell. Transp. Syst.*, vol. 17, no. 9, pp. 2637–2647, Sep. 2016.

[5] A. Bagula, L. Castelli, and M. Zennaro, "On the design of smart parking networks in the smart cities: An optimal sensor placement model," *Sensors*, vol. 15, no. 7, pp. 15443–15467, 2015.

[6] W. Xiao, B. Vallet, K. Schindler, and N. Paparoditis, "Street-side vehicle detection, classification and change detection using mobile laser scanning data," *ISPRS J. Photogramm. Remote Sens.*, vol. 114, pp. 166–178, Apr. 2016.

[7] L. Zhang, X. Li, J. Huang, Y. Shen, and D. Wang, "Vision-based parking-slot detection: A benchmark and a learning-based approach," *Symmetry*, vol. 10, no. 3, p. 64, 2018.

[8] Z. Qian and R. Rajagopal, "Optimal dynamic pricing for morning commute parking with occupancy information," *Transportmetrica A, Transp. Sci.*, vol. 11, no. 4, pp. 291–316, 2015.

[9] O. Zoeter, C. Dance, S. Clinchant, and J.-M. Andreoli, "New algorithms for parking demand management and a city-scale deployment," in *Proc. 20th ACM SIGKDD Inter. Conf. Knowl. Discovery Data Mining*, Aug. 2014, pp. 1819–1828.

[10] C. Dowling, T. Fiez, L. Ratliff, and B. Zhang, "Optimizing curbside parking resources subject to congestion constraints," in *Proc. IEEE 56th Annu. Conf. Decis. Control*, Dec. 2017, pp. 5080–5085.

[11] G. Pierce and D. Shoup, "Getting the prices right: An evaluation of pricing parking by demand in san francisco," *J. Amer. Planning Assoc.*, vol. 79, no. 1, pp. 67–81, Jan. 2013.

[12] L. J. Ratliff, C. Dowling, E. Mazumdar, and B. Zhang, "To observe or not to observe: Queuing game framework for urban parking," in *Proc. IEEE 55th Conf. Decis. Control*, Dec. 2016, pp. 5286–5291.

[13] S. Yang and Z. S. Qian, "Turning meter transactions data into occupancy and payment behavioral information for on-street parking," *Transp. Res. C, Emerg. Technol.*, vol. 78, pp. 165–182, May 2017.

[14] T. Lin, H. Rivano, and F. Le Mouël, "A survey of smart parking solutions," *IEEE Trans. Intell. Transp. Syst.*, vol. 18, no. 12, pp. 3229–3253, Dec. 2017.

[15] J. Glasnapp *et al.*, "Understanding dynamic pricing for parking in Los Angeles: Survey and ethnographic results," in *Proc. Int. Conf. HCI Bus.* Cham, Switzerland: Springer, 2014, pp. 316–327.

[16] X. Ma, X. Sun, Y. He, and Y. Chen, "Parking choice behavior investigation: A case study at Beijing Lama temple," *Procedia-Social Behav. Sci.*, vol. 96, pp. 2635–2642, Nov. 2013.

[17] T. Fiez, L. J. Ratliff, C. Dowling, and B. Zhang, "Data driven spatio-temporal modeling of parking demand," in *Proc. Annu. Amer. Control Conf.*, Jun. 2018, pp. 2757–2762.

[18] Statistical Atlas. (2017). *Population by Neighborhood in Seattle*. [Online]. Available: https://statisticalatlas.com/neighborhood/Washington/Seattle/Belltown/Population

[19] K. P. Murphy, *Machine Learning: A Probabilistic Perspective*. Cambridge, MA, USA: MIT Press, 2012.

[20] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *J. Roy. Stat. Soc., B (Methodol.)*, vol. 39, no. 1, pp. 1–22, 1977.

[21] C. F. J. Wu, "On the convergence properties of the EM algorithm," *Ann. Statist.*, vol. 11, no. 1, pp. 95–103, 1983.

[22] G. Schwarz, "Estimating the dimension of a model," *Ann. Statist.*, vol. 6, no. 2, pp. 461–464, 1978.

[23] P. A. P. Moran, "Notes on continuous stochastic phenomena," *Biometrika*, vol. 37, nos. 1–2, pp. 17–23, 1950.

[24] T. Fiez and L. J. Ratliff, "Data-driven spatio-temporal analysis of curbside parking demand: A case-study in seattle," 2017, *arXiv:1712.01263*. [Online]. Available: https://arxiv.org/abs/1712.01263

Tanner Fiez (S'16) received the B.S. degree in ECE from Oregon State University. He is currently pursuing the Ph.D. degree in ECE with the University of Washington, Seattle. His research interests are in online learning and learning in games. He was a recipient of the National Defense Science and Engineering Graduate Fellowship in 2017.

Lillian J. Ratliff (S'08–M'15) received the B.S. degree in mathematics and the B.S. and M.S. degrees in EE from the University of Nevada, Las Vegas, in 2008, 2008, and 2010, respectively, and the Ph.D. degree from the University of California at Berkeley, in 2015. She was a Post-Doctoral Researcher in EECS with the University of California at Berkeley. She is currently an Assistant Professor of ECE with the University of Washington, Seattle. She was a recipient of the NSF Graduate Research Fellowship in 2009, the CISE CRII Award in 2017, and the NSF CAREER Award in 2019.