Data Driven Spatio-Temporal Modeling of Parking Demand

Tanner Fiez, Lillian J. Ratliff, Chase Dowling, Baosen Zhang[1]

***Abstract*— To mitigate the congestion caused by 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. In light of this, we provide an extensive study of the spatio-temporal characteristics of parking demand. This work advances the understanding of where and when to set pricing policies, as well as how to target information and incentives to drivers looking to park. Harnessing data provided by the Seattle Department of Transportation, we develop a Gaussian mixture model based technique to identify zones with similar spatial demand as quantified by spatial autocorrelation. In support of this technique we provide a method based on the repeatability of our Gaussian mixture model to show demand for parking is consistent through time.**

I. Introduction

Developing effective parking policy is challenging considering the diverse needs of a city which must be balanced. Perhaps the most significant consequence of failing to do so adequately is increased congestion on arterials. Indeed, it has been estimated that approximately 30% of traffic in a city is due to vehicles in search of parking [1], [2]. The economic and environmental impacts of this phenomenon have been shown to be significant [3]–[5].

While performance based pricing strategies aimed at combating adverse impacts of parking are widely researched [6]–[9] and explored by cities [10]–[13], there is growing evidence that several factors beyond price drive parking decisions. The results of surveys conducted in Los Angeles and Beijing on drivers looking to park revealed that proximity to an intended destination influenced decisions more than price [14], [15]. Correspondingly, the results of empirical studies on the price elasticity of parking demand support considering control methods beyond price. One of the most expansive studies examined the SFpark Project [10] and found that price elasticities varied significantly with location and time of day, indicating price was not the only factor in decisions [7].

The complexity of many proposed pricing schemes is also problematic. For instance, it took two price adjustments and increased marketing before drivers were aware enough of policies to change their behavior in the SFPark Project [7]. Likewise, the aforementioned survey in Los Angeles confirmed drivers awareness to price is low [14].

Despite new data sources that could potentially support nuanced management strategies, cities generally employ simple static pricing schemes, predominantly owing to the obstacles to carrying out substantial policy changes. While maintaining salient features of the existing policies which make them viable, e.g., being easy to track and understand, the approach to selecting where and when to set them can be greatly improved by exploiting available data streams. Thus in contrast to prior research, we analyze frequently overlooked factors in parking decisions such as location and time of day, to propose methods that can improve traditional policies with straightforward modifications.

Leveraging publicly available data sources, we develop approaches to identify zones and time periods with similar spatial and temporal parking demand respectively—allowing for more effective simple static pricing schemes. Specifically, we show that a Gaussian mixture model (GMM) can be used to identify groups of block-faces which have a high degree of spatial autocorrelation. We supplement the model by providing a method based on the repeatability of the GMM to metricize the consistency of parking demand, and demonstrate through experiments that demand is indeed consistent through time.

In Section II we introduce our data sources, describe our method to estimate demand, and discuss the spatio-temporal characteristics of parking demand. We describe our approach using a GMM to identify zones with similar spatial demand and how we quantify this using spatial autocorrelation in Sections III and IV, respectively. In Section V, we present the results of our analysis using parking data from Seattle, WA and conclude with a discussion in Section VI.

II. Spatio-Temporal Characteristics of Demand

A. Data Sources

We use paid parking transaction data, block-face supply data, and GPS location data of the block-faces from January 1$^{\text{st}}$, 2017–July 30$^{\text{th}}$, 2017 made available to us via the Seattle Department of Transportation (SDOT)[1]. The paid parking transaction data includes both pay-station and pay-by-phone records for each block-face. Paid parking is available Monday–Saturday. The block-face supply data consists of the estimated number of parking spaces for each block-face[2].

*The authors have been supported in part by NSF grants CSN-1646912 and CNS-1634136. T. Fiez was also supported in part by an NDSEG Fellowship and C. Dowling was also supported in part by the Washington Clean Energy Institute.

[1]T. Fiez, L.J. Ratliff, C. Dowling, B. Zhang are with the Department of Electrical Engineering, University of Washington, Seattle, WA 98195, USA. Emails: {fiezt, ratliffl, cdowling, zhangbao}@uw.edu

Code to reproduce work available at github.com/fiezt/spatial-data-analysis

[1]These data sources are available via the open data portal at data.seattle.gov.

[2]In Seattle, parking spaces are not marked and thus the number of spaces for each block-face is estimated by dividing the length of the legal parking zone into 25 foot increments.

The GPS location data of the block-faces includes the latitude and longitude of both ends of a block-face. We use the endpoints to get the coordinates of the midpoints of block-faces.

B. Demand via Estimated Occupancy

With expanding use of smart parking meters, the most widely applicable method to estimate occupancy is through paid transaction data. In this method, the estimated occupancy at block-face i at time k is given by

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

We estimate the occupancies at each minute and aggregate them to an hour granularity since prices do not change at any higher frequency than this.

The estimated occupancy deviates from the true occupancy because select vehicles are permitted to park for free, vehicles leave before the paid time is up, and the estimated supply of a block-face may be inaccurate due to spaces not being marked. These factors can cause the estimated occupancy to be greater than 100%, and we clip the maximum estimated occupancy at 150%. The estimated occupancy eclipses this limit less than 0.45% of the hourly occupancy instances over all block-faces. Because our analysis focuses on the relative relationship between occupancies, using the estimated occupancy has a negligible affect on our analysis.

C. Testbed

We focus our analysis on the Belltown neighborhood in Seattle. This neighborhood is of particular interest by nature of it being a rapidly growing mixed use development. Moreover, Belltown has both the highest population density [16] and the most complete coverage of on street parking of any neighborhood in Seattle.

Fig. 3. Paid parking divisions in Belltown. Parking in the north zone (red) is \$1.00/hr in 8AM–11AM and \$1.50/hr in 11AM–8PM with four hour time limits. In the south zone (blue) parking is \$2.50/hr in 8AM–5PM and 5PM–8PM with two and three hour time limits respectively.

D. Spatial and Temporal Characteristics

Evaluating the occupancy profiles of Belltown in Fig. 1, we find Monday–Friday see comparable demand, while Saturday follows a different trend. During weekdays, occupancy increases from the opening of paid parking until demand peaks near lunch time, before decreasing during the afternoon until increasing again during the evening hours near dinner time. 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 policies must consider the temporal characteristics of demand that may be driven by businesses.

We use Fig. 2 as a motivating example of the spatial demand characteristics we observe. In Fig. 2a, the occupancy at 7PM on Friday is more or less uniformly distributed throughout Belltown, with the exception of an area of much higher occupancy in the center of the neighborhood. This area happens to have a high concentration of bars and restaurants which we conjecture drive demand. Interestingly, the area of high occupancy also appears to be up against the divide of the north and south paid parking zones—denoted by red and blue block-faces respectively in Fig. 3—which have a $\$1.00$/hr price difference at this time. One could conjecture that a superior division of paid parking zones exists which could reduce the congestion in this area.

In Fig. 2b, the occupancy at 10AM on Saturday has a more diverse distribution, but most importantly the areas of high occupancy are located 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 some of the most 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. This example highlights one of the key problems we seek to address in this paper: parking policies with uniform pricing schemes in arbitrary zones and time periods ignore important spatio-temporal characteristics which reduces their effectiveness.

III. Gaussian Mixture Model

We use a GMM as a clustering method to find zones and groups of block-faces within them that are spatially close and have similar demand. This technique enables us to:

1) Draw new inferences about parking demand and its spatio-temporal characteristics.
2) Use data to make informed decisions about zones and time periods in which static, uniform pricing schemes would be more effective than if chosen arbitrarily.
3) Consider identified zones as groups of users with similar preferences, facilitating targeted information and incentives.

A. Model Description

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

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

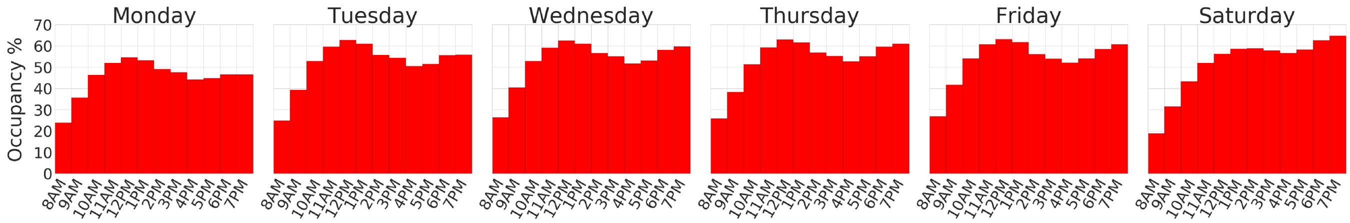

Fig. 1. Mean occupancies at active paid parking times at each day of the week paid parking is available in Belltown.

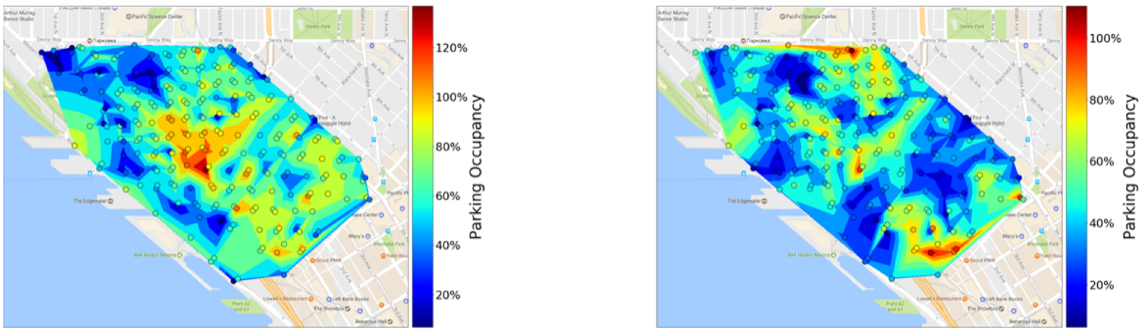

(a) Friday 7PM

(b) Saturday 10AM

Fig. 2. Contours of the mean occupancies in Belltown. Each scatter point is the midpoint of a block-face.

We consider each sample of our dataset to be a vector $x_i \in \mathbb{R}^3$, containing spatial and demand features for a block-face as

$$x_i = \begin{bmatrix} x_{i,\text{latitude}} & x_{i,\text{longitude}} & x_{i,\text{occupancy}} \end{bmatrix}. \tag{3}$$

Thus the complete dataset is given by the matrix of the n samples stacked as $x = \begin{bmatrix} x_1 & \cdots & x_n \end{bmatrix}^T$. In our implementation we normalize features column-wise to be in $[0, 1]$. The motivations for the features we choose is their simplicity—they exactly capture the spatial demand aspects of the data we have—and they work to tradeoff grouping block-faces which are close and block-faces which have similar demand.

We consider a vector of n indicator variables $z = \begin{bmatrix} z_1 & \cdots & z_n \end{bmatrix}$ as the latent component labels for samples. The prior on the probability of a sample belonging to a mixture component can then be expressed as

$$p(z_i = j) = \pi_j. \tag{4}$$

The parameter π must satisfy the restrictions $\pi_j \in [0, 1]$ and $\sum_{j=1}^{k} \pi_j = 1$. The likelihood of a sample belonging to a mixture component is given by

$$p(x_i | z_i = j) = \mathcal{N}(x_i | \mu_j, \Sigma_j), \tag{5}$$

where the multivariate Gaussian distribution is

$$\mathcal{N}(x_i | \mu_j, \Sigma_j) = \frac{\exp\left[-\frac{1}{2}(x_i - \mu_j)^T \Sigma_j^{-1}(x_i - \mu_j)\right]}{(2\pi)^{\frac{3}{2}} |\Sigma_j|^{\frac{1}{2}}}. \tag{6}$$

We make the common assumption that the features are conditionally independent given the component, i.e. each covariance matrix Σ_j is diagonal.

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

$$\text{LL} \triangleq \log\ p(x|\pi, \mu, \Sigma) = \sum_{i=1}^{n} \log \sum_{j=1}^{k} \pi_j \mathcal{N}(x_i | \mu_j, \Sigma_j). \tag{7}$$

We employ the expectation-maximization (EM) algorithm [17] to optimize for this objective. The EM algorithm, given in Algorithm 1, consists of an initialization of the unknown parameters and two steps, the E step and the M step, which are repeated until convergence. The convergence criteria we use is to terminate the algorithm when the change in the log likelihood between iterations, which is ensured to be positive since the log likelihood is guaranteed to increase at each iteration of the EM algorithm [18], [19], is smaller than a parameter ϵ.

$$\Delta\text{LL} \triangleq \text{LL}^i - \text{LL}^{i-1} \leq \epsilon. \tag{8}$$

In the E step, the expected values of the unobserved component labels given the current parameter values are updated. These are the posterior probabilities and are sometimes referred to as the responsibility that component j takes for data point i [20]. Formally, we will denote this term as

$$r_{i,j} \triangleq p(z_i = j | x_i, \pi_j, \mu_j, \Sigma_j). \tag{9}$$

In the M step, the parameter values are updated to maximize the log likelihood.

Once the convergence criteria is met, we make hard assignments of each sample x_i to the component label j which maximizes the responsibility $r_{i,j}$—that is,

$$z_i^* = \operatorname*{argmax}_j\ r_{i,j}. \tag{10}$$

The objective function is non-convex, which only guarantees that we find local minima. Hence, we run the algorithm

Algorithm 1 EM Algorithm for GMM

procedure EM(x) $\triangleright$ x: normalized feature matrix
  for each initialization **do**
    while $\Delta\text{LL} > \epsilon$ **do**
      for each sample x_i **do**
        E-Step:
        for j in $\{1, \ldots, k\}$ **do**

$$r_{i,j} = \frac{\pi_j \mathcal{N}(x_i|\mu_j, \Sigma_j)}{\sum_{j'}^{k} \pi_{j'} \mathcal{N}(x_i|\mu_{j'}, \Sigma_{j'})}$$

        end for
      end for
      for each component π_j, μ_j, Σ_j **do**
        M-Step:

$$\pi_j = \frac{1}{n} \sum_{i=1}^{n} r_{i,j}$$

$$\mu_j = \frac{\sum_{i=1}^{n} x_i r_{i,j}}{\sum_{i=1}^{n} r_{i,j}}$$

$$\Sigma_j = \frac{\sum_{i=1}^{n} (x_i - \mu_j)^T (x_i - \mu_j) r_{i,j}}{\sum_{i=1}^{n} r_{i,j}}$$

      end for
    end while
    Store maximized LL for the initialization
    Store sample assignments using (10)
  end for
end procedure

for several random initializations and retain the model from the iteration that resulted in the highest log likelihood.

The GMM also involves a model selection problem of selecting the number of mixture components. We leverage the Bayesian Information Criterion (BIC) [21] to solve this problem. In particular, we select the number of components which minimizes the BIC.

B. Consistency Metric

We aim to quantify how similar demand is from week to week at a given day of week and time of day, i.e. we want to determine the *consistency* of demand. Pricing schemes, as well as targeted information and incentive campaigns, can be constructed more effectively knowing that demand characteristics will remain the same without changes to policy or the system.

We propose a method to metricize the consistency of demand based on the repeatability of our GMM approach. Using our dataset, the procedure to determine the consistency metric value at a day of the week and hour of the day is as follows:

1) For the chosen day of the week and hour of the day, select a specific date and fit a GMM using the occupancy data at this instance.
2) Assign component labels to each block-face for all other instances with the same day of the week and hour of the day in the dataset using the learned model.
3) Determine the percentage of block-faces which were assigned to the same component as they were in the original GMM fit.
4) Repeat (1)–(3) switching the date on which the GMM is fit, and then average over the percentages computed at each iteration.

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

IV. Spatial Autocorrelation

We use a standard measure of spatial autocorrelation—Moran's I [22]—to quantify the degree of spatial homogeneity or heterogeneity present in the demand. Moran's I is defined as

$$I = \frac{N}{\sum_i \sum_j w_{i,j}} \frac{\sum_i \sum_j w_{ij}(o_i - \bar{o})(o_j - \bar{o})}{\sum_i (o_i - \bar{o})^2}, \qquad (11)$$

where for our problem 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 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.

There are several ways the spatial weight matrix W can be designed depending on the objective. We explore three such methods to evaluate certain questions of interest in Section V. In particular, we will evaluate each method 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 in our data set—each instance given by the occupancy at a date and time—that are significant.

A. Assessing Local Homogeneity

In Section II we discussed and demonstrated that parking demand displays spatial heterogeneity. A logical follow up question to this observation is whether there is at least local spatial homogeneity. If this were the case, it would imply that it could be possible to find groups of block-faces where there is spatial homogeneity. To evaluate this objective we create the weight matrix by setting values of w_{ij} to 1 if block-face j is one of the k nearest neighbors to block-face i and 0 if it is not. We experiment using a range of values for k in Section V.

B. Assessing Homogeneity in Current Parking 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 appraise 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 1 if block-faces i and j are in the same parking zone and 0 if they are not.

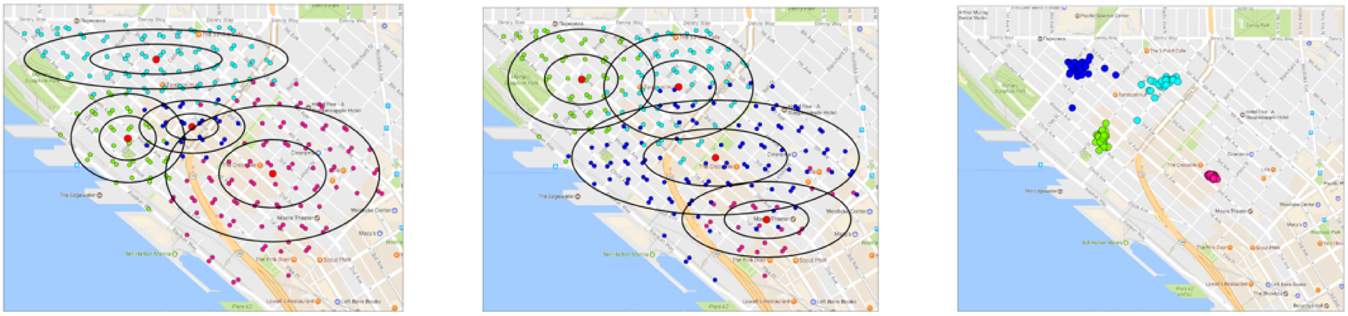

(a) Friday 7PM (b) Saturday 10AM (c) Wednesday 10AM

Fig. 4. GMMs fit using mean occupancies in Belltown at Friday 7PM in Fig. 4a and Saturday 10AM in Fig. 4b. Block-faces are colored by the component label, the ellipses indicate the first standard deviations of the components GPS coordinates, and the red scatter points indicate the centers of the components GPS coordinates. Fig. 4c shows clusters of the spatial centers of the components from GMM fits in Belltown on all Wednesdays in our dataset at 10AM.

TABLE I

CONSISTENCY METRIC RESULTS OVER ALL PAID DAYS OF THE WEEK AND HOURS OF THE DAY. UNITS ARE IN PERCENTAGE.

Day\Time	8AM	9AM	10AM	11AM	12PM	1PM	2PM	3PM	4PM	5PM	6PM	7PM	**Daily Consistency**
Monday	67.7	78.1	83.5	87.9	90.9	84.5	84.5	82.1	82.6	83.3	78.7	79.5	**81.9**
Tuesday	69.3	86.0	90.7	92.7	91.6	90.0	89.7	87.4	83.1	85.4	85.8	85.9	**86.5**
Wednesday	72.9	87.0	90.9	91.5	91.4	90.1	90.6	86.5	85.6	86.3	84.3	85.6	**86.9**
Thursday	71.0	84.3	88.6	91.7	90.1	90.6	90.3	88.1	86.4	87.2	84.8	86.8	**86.7**
Friday	72.2	86.0	89.8	90.9	89.6	90.1	90.6	87.9	87.1	87.0	84.5	85.4	**86.7**
Saturday	63.3	74.8	79.7	83.8	86.4	84.6	85.2	84.9	83.6	85.7	86.5	83.0	**81.8**
Hourly Consistency	**69.4**	**82.7**	**87.2**	**89.7**	**90.0**	**88.3**	**88.5**	**86.2**	**84.7**	**85.8**	**84.1**	**84.3**	

C. Assessing Homogeneity in GMM Components

One of the aims of the GMM approach is to identify groups of block-faces that are spatially close and have similar demand. To gauge our success in doing so, and to justify considering the zones as groups of users with similar preferences, we create the weight matrix by setting values of w_{ij} to 1 when block-faces i and j are in the same GMM component and 0 when they are not.

V. EXPERIMENTS & RESULTS

We now explore the application of our GMM approach, provide analysis of the spatio-temporal characteristics and consistency of parking demand, and demonstrate the advantages of our approach by considering spatial autocorrelation.

A. Modeling Belltown with GMM

In Fig. 2, we illustrate the mean spatial demand in Belltown at Friday 7PM and Saturday 10AM. Figs. 4a and 4b provide an example use of our GMM approach with the same data. It is clear that we are able to find separable zones in which spatially close block-faces are included in the same mixture components. This is important due to the fact that while there may be spatial heterogeneity in Belltown, we are able to find zones in which block-faces have similar demand thereby validating that zone based pricing is viable.

Indeed, we find that within Belltown there is local homogeneity which enables our GMM approach. Using the method described in the previous section to determine the spatial homogeneity locally, we find that the spatial autocorrelation is significant 91.9%, 96.67%, and, 98.2% of the time using 3, 5, and 10 nearest neighbors, respectively. Leveraging this, we are able to fit a GMM where the spatial autocorrelation is significant 99.9% of the time in our dataset. This is a major improvement over the current paid areas as they only have significant autocorrelation in 66.4% of the instances.

The example depicted in Figs. 4a and 4b also indicates that the model we learn is related to the day of the week and time of day. The model we learn for Friday night, e.g., is very different from the model we learn for Saturday morning, asserting that the spatial component of demand depends on the temporal component. Consequently, in the design of pricing policies and information schemes, the questions of where and when to designate them should be considered together.

An interesting artifact in our analysis of Belltown is the model selection problem. Using the model selection criterion described in Section III-A, we find four mixture components—corresponding to four paid parking zones—to be optimal. At present, Belltown has just two paid parking zones in place. This may play a significant factor in why we can improve on the existing policy design method of using heuristics to set paid zone boundaries.

B. Consistency of Parking Demand

The results in Table I establish that parking demand is typically consistent through time. With the exception of the first hour of paid parking in a day when occupancy is very low as drivers arrive, the average consistency value at an hour of a day is very high ranging from 82.7%–

90.0%. Furthermore, on a given day of the week, even when including the less consistent first hour of the day, the average consistency value is still very high ranging from 81.8%–86.9%.

We also investigate how the spatial centers of the mixture components change from week to week. We do this for a day of the week and hour of the day by using the k-means clustering algorithm [20] on the centers of components that were found at each date with the same corresponding day of the week and hour of the day. Fig. 4c shows an example of clustering the centers from 29 GMM fits on different Wednesdays at 10AM. By finding the centroids of each of the k-means clusters and calculating the average distance from each centroid to the points in that respective cluster, we can describe this change in terms of distance. In Fig. 4c, we find the average distance of the points to their respective centroids to be just 29m. The corresponding value at all other days of week and hours of the day has mean of 69.3m. Thus we can see the spatial centers are reliable in our model across time to within just a few street blocks.

C. Spatio-Temporal Insights

A key insight we gain from the GMM approach is learning more about the time periods in which spatial demand is similar. We find that Monday–Friday from 8AM–4PM nearly identical models are learned. Likewise, for Monday–Friday from 4PM–8PM very similar models are learned, which are different from those learned Monday–Friday from 8AM–4PM. Contrarily, models we learn for Saturday are quite unique and need to be considered on their own[3].

The preceding observations indicate that, based off of our model, it would make the most sense to have two weekday pricing periods—i.e. 8AM–4PM and 4PM–8PM—for the zones we commonly find at these respective time periods, and a unique Saturday pricing scheme. These results are compelling because they are quite different than the policies in place now. Currently, the pricing periods in Belltown are from 8AM–11AM and 11AM–8PM with no individual consideration given to Saturdays.

VI. Discussion & Future Work

We provide an in depth analysis of the spatio-temporal characteristics of parking demand using real data, as well as an interpretable way to find zones where there is spatial homogeneity using a GMM. The work has the potential to allow for more informed decision-making in both policy decisions and in designing targeted information and incentive schemes. We establish that parking demand is consistent, which is to say that without changes to management or infrastructure, learned models will hold up over time. While we focus on Seattle, our methods leverage a now common data source of paid parking transactions from cities, making the models and analysis we use flexible enough to be applied in many other communities.

We seek to use this work to identify zones of similar demand in support of designing targeted information and incentive schemes. Towards this end, we are investigating a multi-arm bandit framework to learn user responses while matching information and incentives in GMM identified zones. We plan to implement developed strategies in a living lab setting in Seattle neighborhoods with the aid of SDOT.

[3] We make an animation available showing the model learned at each day of the week and hour of the day at `github.com/fiezt/spatial-data-analysis/blob/master/animation/mixture.mp4`.

References

[1] D. Shoup and H. Campbell, "Gone parkin'," *The New York Times*, vol. 29, 2007.

[2] C. Dowling, T. Fiez, L. Ratliff, and B. Zhang, "How much urban traffic is searching for parking?" *arXiv preprint arXiv:1702.06156*, 2017.

[3] INRIX. (2017) Searching for parking costs americans $73 billion a year. [Online]. Available: http://inrix.com/press-releases/parking-pain-us/

[4] J. I. Levy, J. J. Buonocore, and K. Von Stackelberg, "Evaluation of the public health impacts of traffic congestion: a health risk assessment," *Environmental health*, vol. 9, no. 1, p. 65, 2010.

[5] K. Zhang and S. Batterman, "Air pollution and health risks due to vehicle traffic," *Science of the total Environment*, vol. 450, pp. 307–316, 2013.

[6] Z. Qian and R. Rajagopal, "Optimal dynamic pricing for morning commute parking with occupancy information," *Transportation Research Part B*, 2012.

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

[8] 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. Knowledge Discovery and Data Mining*. ACM, 2014, pp. 1819–1828.

[9] C. Dowling, T. Fiez, L. Ratliff, and B. Zhang, "Optimizing curbside parking resources subject to congestion constraints," in *Proc. 56th IEEE Conf. Decision and Control (arXiv:1703.07802)*, 2017.

[10] SFMTA. (2017) Sfpark. [Online]. Available: http://sfpark.org

[11] SDOT. (2017) Paid parking. [Online]. Available: http://www.seattle.gov/transportation/parking/paidparkingupdates.htm

[12] LADOT. (2017) La express park. [Online]. Available: http://www.laexpresspark.org

[13] DDOT. (2017) Parkdc. [Online]. Available: https://ddot.dc.gov/page/parkdc

[14] J. Glasnapp, H. Du, C. Dance, S. Clinchant, A. Pudlin, D. Mitchell, and O. Zoeter, "Understanding dynamic pricing for parking in los angeles: Survey and ethnographic results," in *Inter. Conf. HCI in Business*. Springer, 2014, pp. 316–327.

[15] X. Ma, X. Sun, Y. He, and Y. Chen, "Parking choice behavior investigation: A case study at beijing lama temple," *Procedia-Social and Behavioral Sciences*, vol. 96, pp. 2635–2642, 2013.

[16] StatisticalAtlas. (2017) Population by neighborhood in seattle. [Online]. Available: https://statisticalatlas.com/neighborhood/Washington/Seattle/Belltown/Population

[17] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the em algorithm," *J. Royal Statistical Society. Series B (Methodological)*, pp. 1–38, 1977.

[18] C. J. Wu, "On the convergence properties of the em algorithm," *The Annals of statistics*, pp. 95–103, 1983.

[19] R. A. Boyles, "On the convergence of the em algorithm," *J. Royal Statistical Society. Series B (Methodological)*, pp. 47–50, 1983.

[20] K. P. Murphy, *Machine Learning: A Probabilistic Perspective*. The MIT Press, 2012.

[21] G. Schwarz *et al.*, "Estimating the dimension of a model," *The annals of statistics*, vol. 6, no. 2, pp. 461–464, 1978.

[22] P. A. Moran, "Notes on continuous stochastic phenomena," *Biometrika*, vol. 37, no. 1/2, pp. 17–23, 1950.