Applied Geography 89 (2017) 133–141

Contents lists available at ScienceDirect

Applied Geography

journal homepage: www.elsevier.com/locate/apgeog

Predicting demand for 311 non-emergency municipal services: An adaptive space-time kernel approach

Li Xu, Mei-Po Kwan*, Sara McLafferty, Shaowen Wang

Department of Geography and Geographic Information Science, University of Illinois at Urbana-Champaign, Natural History Building, MC-150, 1301 W Green Street, Urbana, IL 61801, USA

ARTICLE INFO

Keywords:
311
Spatial point process
Modeling
Space-time kernel estimation

ABSTRACT

Many cities in the United States and Canada offer a 311 helpline to their residents for submitting requests for non-emergency municipal services. By dialing 311, urban residents can report a range of public issues that require governmental attention, including potholes, graffito, sanitation complaints, and tree debris. The demand for these municipal services fluctuates greatly with time and location, which poses multiple challenges to effective deployment of limited resources. To address these challenges, this study uses a locally adaptive space-time kernel approach to model 311 requests as an inhomogeneous Poisson process and presents an analytical framework to generate predictions of 311 demand in space and time. The predictions can be used to optimally allocate resources and staff, reduce response time, and allow long-term dynamic planning. We use a bivariate spatial kernel to identify the spatial structure and weigh each kernel by corresponding past observations to capture the temporal dynamics. Short-term serial dependency and weekly temporality are modeled through the temporal weights, which are adaptive to local community areas. We also transform the computation-intensive parameter estimation procedure to a low dimensional optimization problem by fitting to the autocorrelation function of historical requests. The presented method is demonstrated and validated with sanitation service requests in Chicago. The results indicate that it performs better than common industry practice and conventional spatial models with a comparable computational cost.

1. Introduction

311 is a special telephone number supported in many cities in the United States and Canada. When the first 311 call center was launched in Baltimore, Maryland in October 1996 as a police non-emergency number, the service was intended to divert non-emergency community concerns from 911 which is reserved for emergency service. 311 calls are assigned a lower priority to make sure that true emergency calls are answered and addressed with the highest priority. However, 311 now has become an important standalone telephone helpline that provides access to a variety of non-emergency municipal services. By dialing 311, urban residents can request services for a range of public issues that require governmental attention, such as graffiti removal, pothole repair, sanitation code violations, and abandoned vehicles. These services are generally implemented at the local level by related government departments and agents, while the available services could vary across cities.

Currently there are over 70 cities and counties in the United States and 18 cities in Canada offering 311 non-emergency services ('3-1-1' 2016). Some European countries have also introduced similar non-emergency helplines, including 11414 in Sweden, 112 in Finland, 115 in Germany, and 101 in the United Kingdom. Most of the cities and counties that offer such helplines have experienced a dramatic increase in the demand for these municipal services since their launch. For example, the Chicago 311 call center receives approximately 3.9 million calls annually in recent years (City of Chicago, 2016), making it an indispensable service tool for the city's residents. The request volume can also become extremely high due to unexpected situations such as major social events. On December 20, 2005, for instance, the 311 call center in New York City received 240,000 calls because it was the first day of the 2005 New York City transit strike ('3-1-1' 2016). To address the rapidly expanding demand for service, Chicago initiated the first comprehensive 311 system to record, track and efficiently document all requests for non-emergency municipal services from intake to resolution. Now many major cities in the United States have adopted an open standard platform, Open311, through which citizens can create and check the status of their service requests online via the Internet. The system now also allows citizens to submit requests directly from a

* Corresponding author.
E-mail addresses: lixu3@illinois.edu (L. Xu), mpk654@gmail.com (M.-P. Kwan), smclaff@illinois.edu (S. McLafferty), shaowen@illinois.edu (S. Wang).

https://doi.org/10.1016/j.apgeog.2017.10.012
Received 30 April 2017; Received in revised form 29 October 2017; Accepted 29 October 2017
0143-6228/

smartphone app.

As argued in many emergency service demand modeling studies, the ability to predict demand is of paramount importance (Goldberg, 2004). This is also true for non-emergency municipal services. Besides the increasing volume of 311 requests, the demand for these municipal services fluctuates a lot with time and location. Calls that arrive at a call center are first evaluated, including obtaining an address, determining the nature and importance of the request, and possibly providing some instructions to the caller (Channouf, L'Ecuyer, Ingolfsson, & Avramidis, 2007). The next step is to check the availability of crew and schedule the dispatch of a service vehicle. When scheduled crew arrive at the scene of call, they assess the situation, provide on-site service, and determine if additional help is necessary. After the completion of service, the crew may travel to the next scene or report to the call center and become available to take new requests. Due to the large amount and great spatial variation of these requests in short time periods, the resources and staff are always limited for providing effective services. As a result, 311 call centers face challenges to maintain a good balance between high service quality and efficiency at multiple temporal scales. First, there is the problem of resource acquisition (Aksin, Armony, & Mehrotra 2007). Long-term forecasts of future incoming call volumes are needed for the strategic planning of staffing levels, system expansion and reorganization. Then, medium-term level forecasts are crucial to make scheduling decisions for 311 planners to dispatch crew and vehicles. This is the problem of resource deployment (Avramidis et al., 2010). Finally, decisions could be informed by updated short-term forecasts in order to adapt to unexpected issues that involve rescheduling and fluctuations in demand, such as those associated with major social events. To address these challenges with effective scheduling and routing design, it is crucial to develop reliable and accurate demand forecasts of future 311 requests in space and time.

Analyses of 311 data have been largely motivated by the need to inform higher-level planning and budget recommendations, and support citizen relationship management such as inspiring citizens to become more involved in neighborhood improvement (International City/County Management Association, 2009; Schellong, 2008). A few other studies employed 311 reports to explore custodianship and expression of territoriality in urban neighborhoods (O'Brien, 2016, 2015; O'Brien, Gordon, & Bladwin-Philippi, 2014) and to study social disparities by race, education, and income (Clark, Brudney, & Jang 2013). Despite the significant utility of 311 services to urban residents, few efforts have addressed the need to accurately predict the spatial and temporal variations of demand in this non-emergency sector. In addition to the complex spatial and temporal patterns in 311 requests, there is a problem of data sparsity. Since the requests are highly unevenly distributed in space, zero observations appear at many locations, especially at fine spatial and temporal granularities.

Although similar needs exist for other service sectors, such as ambulance requests, the industry practice relies on a simplistic local averaging approach. In general, forecasting has been based on simple averaging of historical demand data assuming a similar future trend within a geographic region (Goldberg, 2004). For example, the Toronto Emergency Medical Service predicts demand by averaging four historical counts in the corresponding hour of the year, from the same geographic region and over the preceding four years (Goldberg, 2004). Similarly, the Charlotte-Mecklenburg (North Carolina) Emergency Medical Service agency (MEDIC) makes predictions by averaging twenty historical counts in the corresponding hour of the past four weeks over the previous five years (Setzler, Saydam, & Park, 2009). These simplistic approaches do not take into account the complex spatial and temporal variations in demand for service and are thus prone to yielding inaccurate and noisy predictions. Further, these methods are sensitive to the zoning of spatial units as they average historical counts in arbitrarily defined spatial cells (e.g., 1-km by 1-km squares in Toronto and 4-mile by 4-mile by MEDIC). They thus tend to cause inefficient resource deployment and face the modifiable areal unit problem (MAUP) (Openshaw, 1984).

Early demand modeling studies in the emergency service sector focused on various regression models that incorporate socioeconomic variables (Aldrich, Hisserich, & Lave, 1971; Cadigan & Bugarin, 1989; Kamenetzky, Shuman, & Wolfe, 1982), which were able to generate good predictions of the total demand in a large region, such as a city or a county. However, forecasts are needed at a much finer spatial and temporal granularity for effective deployment purpose. Trudeau, RousseauFerland, and Choquette (1989) recognized the spatial and temporal patterns in demand for emergency medical services and applied a logit market share model to produce forecast in 3-hour time blocks. However, their model did not incorporate spatial variations in making demand predictions. Channouf et al. (2007) compared various time series models and developed autoregressive moving average models that can make accurate daily forecasts into the future. They suggested the inclusion of the spatial patterns of demand based on time and demographic characteristics in future studies. More temporal process based approaches were proposed later, including time series with dynamic latent factor structure (Matteson, McLean, Woodard, & Henderson, 2011) and singular spectrum analysis (Vile, Gillard, Harper, & Knight, 2012). Setzler et al. (2009) designed an artificial neural network (ANN) model to forecast at various spatial and temporal resolutions simultaneously. However, ANN models were not statistically better than the simplistic method used by MEDIC in predicting demand at the finest spatial granularity, partially due to the data sparsity problem. Zhou et al. (2014) developed a time-varying Gaussian mixture model to predict ambulance demand for Toronto, Canada, which was proven to be more accurate than industry practice. However, the model requires expert experience as prior knowledge to make good predictions, and the parameter estimation procedure is highly computationally intensive.

Kernel density estimation (KDE) is a nonparametric method for estimating a smooth probability density function in statistics (Silverman, 1986). It is now also a spatial analysis technique for converting point or line features into density surfaces in GIS (e.g., Kwan, 2000, 2004). Brunsdon, Corcoran, and Higgs (2007) evaluated various visualization techniques, including map animation, the comap (Brunsdon, 2001) and the isosurface, and extended KDE to investigating spatiotemporal phenomena. KDE and scan statistics based approaches were used to model and forecast the spatiotemporal patterns of a wide range of data, including ambulance requests (Zhou & Matteson, 2015), disease risk (Zhang et al., 2011), and crime clusters (Brunsdon et al. 2007; Nakaya & Yano, 2010). Most of these studies create an independent spatial kernel for each time block or combine a spatial kernel with a temporal kernel without taking into account interactions between the spatial and temporal dimensions. However, Schoenberg (2004) argued that most space-time processes are not completely separable, and further described and compared two nonparametric tests to examine the separability of a space-time marked point process. Without assuming complete separability, Mohler et al. (2011) applied a self-exciting point process commonly used in seismology to urban crimes. Their method combines stochastic declustering and KDE to model the space-time triggering function and temporal trends. Temporally weighted KDE method have been used for next event prediction, which weighs the conditional spatial density function more by recent events (Bowers, Johnson, & Pease 2004; Johnson, Bowers, Birks, & Pease, 2009; Porter & Reich, 2012). Zhou and Matteson (2015) proposed an alternative that combines a conditional temporal weight function with a spatial kernel. No study to date has attempted to improve demand predictions of 311 service requests by leveraging their spatiotemporal characteristics.

To be able to better respond to 311 service requests, we need a robust method that can accurately predict these requests at reasonably high spatial and temporal resolution. To meet this challenge, we propose a locally adaptive space-time kernel estimation model that utilizes the spatiotemporal characteristics of 311 service requests to enhance predictive accuracy. The model is adapted from the stKDE approach