Correlational Analysis between Weather and 311 Service Request Volume

Jordan Ponn, Mark S. Fox

Enterprise Integration Laboratory (eil.utoronto.ca)
Mechanical and Industrial Engineering
University of Toronto
5 King's College Road, Toronto ON, M5S 3G8
2 May 2017

Jordan Ponn: j.ponn@mail.utoronto.ca
Mark S. Fox: msf@eil.utoronto.ca

Abstract

This article investigated the relationship between 311 Toronto service requests and weather phenomena in order to develop a model to predict service request call volume. Daily service request volumes, grouped by FSA, are compared against weather events and weather attributes, such as rainfall, temperature and snow. Correlations were found to be weak to moderate in strength, which was not sufficient for building a predictive model.

1. Introduction

311 is a North American municipal hotline used to handle non-emergency service requests. For residents, it acts as the first point of contact with the City for information about City services and programs. It was first created by the White House and the Office of Community Oriented Policing Services (COPS) of the Department of Justice to alleviate call volume congestion from the 911 system for emergency services. Originally piloted in Baltimore, Maryland in 1996, it has since spread to other cities across North America. While initially formed to give police officers more opportunities to be engaged with the community, the system has grown to consolidate numerous city services (Schwester, Carrizales, & Holzer, 2009).

In the summer of 2009, Toronto launched its own 311 service. It is the largest service of its kind in Canada that provides information and services 24/7. Since its inception, it has consolidated 251 inquiry lines across 9 call centers, and can provide services in over 180 languages. From 2010 to 2014, there has been an average yearly 5.5% growth in call volume, where in 2014, they received 1.18 million calls. As of 2016, it employs 101.5 full time Customer Service Representatives and 8 coordinators to support the system[1].

[1] http://www.toronto.ca/311

The service accepts requests from a variety of mediums. While phone calls are the primary method of contact, requests can also be made through Email, Twitter, as well as through the 311 website. In 2012, 311 Toronto launched "Mobile Apps", an API which allowed the public to submit pothole and graffiti service requests via Mobile[1].

Call centers are constantly under pressure to reduce costs and improve quality of customer interactions, and 311 Toronto is no exception - operation costs for 2016 was budgeted for $17.301 million gross and $8.982 million net (after adjusting for revenues from service delivery, as well as business and information processing), with 74% of expenses allocated towards service delivery[2]. Being able to forecast service request volume would allow for better staffing schedules to meet the anticipated demand, and provide a better customer experience. In addition, analyzing these service requests can provide a clearer picture on the health of the city. This can enable data-driven decisions to better allocate city resources to deal with issues before they become problems.

The remainder of this paper is as follows. Section two reviews some of the background work that has gone into forecasting call center volumes. The third section will review the data sets that are used, and the fourth section will go over the analysis methodology. The fifth section will cover the characteristics of the data, followed by the analyses. The sixth section will go over the results and seventh will discuss future work.

2. Background

When modeling call center volume, the quantity of calls are aggregated over a specified interval time, and forecasting is done to predict this volume at later times. In total, there are two types of forecasts; inter-day, which predicts days in advance, and intra-day, for predictions within a given day.

Initially, inter-day call volume forecasting was the primary prediction method due to the lack of quality data. Using an autoregressive integrated moving average (ARIMA) on aggregated call counts, Andrews and Cunningham noticed the how holidays and marketing affected call volume (Andrews & Cunningham, 1995). Jongbloed and Koole proposed using a Poisson process to model forecasts (Jongbloed & Koole, 2001), while Brown et al. (Brown, et al., 2005) and Avramidis et al. (Avramidis, Deslauriers, & L'Ecuyer, 2004) identified call volume intra-day dependence.

Weinberg expanded on their work by using a Bayesian framework with Poisson arrival rates to feed a Monte Carlo Markov chain model (Weinberg, Brown, & Stroud, 2007). This allowed for inter-day and intra-day forecasts, however it was computationally expensive to calculate. Shen and Huang significantly improved upon this by using singular value decomposition to reduce dimensionality, and split factors into inter- and intra- day predictors (Shen & Huang, 2008). Taylor performed a study which compared an array of methods, and showed that the effectiveness of any given methodology was affected by lead times and workloads (Taylor, 2008).

[2] http://www.toronto.ca/budget2016

Other researchers have also tried using regression methods to predict call volumes. Klungle and Maluchnik used a dynamic regression model to predict class for emergency road services (Kungle & Maluchnik, 1998), while Ibrahim and L'Ecuyer tried using a linear regression model with independent residuals (Ibrahim & L'Ecuyer, 2012). Pacheco et al. looked at using a backpropagation neural network (Pacheco, Millán–Ruiz, & Vélez, 2009), while Millan-Ruiz and Hidalgo expanded on a neural network-based model for intraday forecasts (Millán–Ruiz & Hidalgo, 2013).

311 Environment

311 Toronto receives calls can be classified into two categories – general inquiries and actionable requests. As actionable requests are a result of a complaint or a perceived problem that has occurred in the city, it is reasonable to believe that some service requests would be influenced by additional factors. For example, excessive rains may cause flooding related problems, while festivals may cause in increase in noise and litter complaints.

Zha and Veloso did a study on the New York City 311 service, and built a predictive model using a random forest trained with 3 years of service request data (Zha & Veloso, 2014). In their model, they factored in day of the week, holidays, mean temperatures, snow, as well as the last 7 days of call volumes. They identified that the strongest predictor for a given day's call volume was to use the counts of the previous 7 days, and that their weather inputs had little influence over call volumes.

As a continuation to their study, we are looking to see if integrating more detailed weather data into the model will act as a stronger predictor.

3. Data Profile

As part of the Open Government movement to make city information more accessible to the general public, the City of Toronto publishes and maintains over 200 of machine readable datasets in the Toronto Open Data Catalog[3]. This repository was sourced for indicators related to 311 Toronto activities. Statistics Canada[4] and Environment Canada[5] were also sourced respectively for population and historical weather datasets.

311 Toronto Service Requests

Records of customer initiated service requests received by 311 Toronto from January 2010 to June 2015 are publically available from the Toronto Open Data Catalog[3]. Combined, the datasets are comprised of 1.83 million records with 3 attributes, where each record is an individual service request. The three attributes are as follows:

[3] http://www.toronto.ca/open
[4] http://www.statcan.gc.ca
[5] http://climate.weather.gc.ca

Creation Date: Date time that the service request was created.

Location: Description of where the service request happened. The description is listed as either the Forward Section Area (FSA) code, which is the first three characters of a 6-digit postal code, or as an intersection description in the format: [*Street Intersection*] || [*District Name*].

Service Request Type: Description of the service request that was created.

Also available on the Toronto Open Data Catalog[3] is a service request codes dataset, which describes the different types of service request problems. This dataset lists 371 service request types available as of November 2010. Each record has 5 attributes:

Division: Department responsible for handling the request (i.e. Solid Waste Management)

Section - Unit: Group within a division responsible for handling the service request

Service Request Name: Name of the service request.

Problem Code: Identification code of the service request

Internet Self Serve: Indicator if the service request can be created online

There is over 45 service divisions, but only 5 that 311 Toronto will initiate requests for, which are represented in the datasets. These divisions are: *Municipal Licensing & Standards*, *Solid Waste Management Services*, *Toronto Water*, *Transportation Services*, and *Urban Forestry (Park, Forestry & Recreation)*. These divisions can be further broken down into section units which are subgrouping of individual service request types. Examples include *District Enforcement* under *Municipal Licensing & Standards* or *Road Operations* under *Transportation Services*.

Geopolitical Boundaries and Populations

Population per FSA was sourced from Statistics Canada[4] from the 2011 Census, and was used to normalize service request counts in the analysis. Each record contained 4 attributes:

Geographic Name: Geographic area, represented by a three character FSA code

Population, 2011: Total population in the given area

Total private dwellings, 2011: Total number of living spaces with a private entrance

Private dwellings occupied by usual residents, 2011: Number of private dwellings occupied by permanent residents

Shapefiles, geographic boundary outlines of each FSA, were also available from Statistics Canada[4].

Weather and Weather Events

Historical weather data is publically available from Environment Canada[5]. Daily weather data from January 2010 to June 2015 was collected from the Pearson Airport, Dustan Airport, North York, and Toronto City weather stations. Daily weather data contained attributes for *average*, *maximum*, and *minimum temperatures*, as well as *precipitation* breakdowns between *snow* and *rain*, and *amount of snow on the ground*.

An hourly weather dataset from Pearson Airport was also retrieved, which contained *temperature*, *relative humidity*, *wind speeds*, and qualitative descriptors for that hour's *weather condition*. Example conditions include '*Cloudy*', '*Clear*', and '*Snow*'. Multiple descriptors are used where applicable (i.e. '*Heavy Rain, Fog*'.

Heat Warning events were sourced from the Toronto Open Data Catalog[1]. A heat alert is issued when the following forecast conditions occur over two consecutive days between May 15th and September 30th:

- Temperature high equal or greater than 31°C
- Temperature low equal or greater than 20°C
- Humidex equal or greater than 40

An extreme heat alert follows the same forecast conditions over the period of 3 days.

This dataset contained 178 records, each with three attributes, where each record is a heat alert or extreme heat alert warning from June 2001 to Sept 2015. The three attributes are as follows:

Date: Date that the alert was issued

Code: A 2-4 character code identifying the type of alert issued

Text: Description of the alert that was issued

4. Methodology

Two analyses were performed to see the effect of weather on service requests – one for inter-day, and one for intra-day. The intra-day analysis evaluated if there was a noticeable change in request volume on a day when a given weather event occurred, while the inter-day analysis evaluated the correlations between daily request volumes to quantifiable daily weather conditions.

For the weather event analysis, hourly weather data was matched with service requests based on date and hour. Of all the available weather stations, the station at Pearson Airport was the only one that had hourly data detailing weather conditions, precipitation amounts, and temperatures. To control for variations in precipitation and temperature across the city, only services requests from the closest Toronto district, Etobicoke, was used in this analysis. Service request counts were grouped based on weather event observed. If a day observed multiple weather events, the daily count was included in each event. A day with only clear or cloudy events was counted towards a 'No Weather Event' category to be used as a baseline. The average baseline daily requests were compared to average daily volumes of days with observed weather events, such as rain or heat alerts, to see the effect of a given weather event on service request volume.

For the correlation analysis, service requests were counted on a daily bases per FSA code and by request type or group. Holidays were then identified based on date using a calendar. Daily request counts for each FSA were compared against each weather feature via Pearson correlation to see the influence of weather on request volume. Weather data was merged to the request count dataset based on date and FSA. Weather and FSA were matched based on the proximity of the center of the FSA to the geo-coordinates of the closest weather station.

The following sections go into more detail as to how the data sets were prepared.

Dates and Holidays

Categorical date features, such as month or day of the week, were extracted based on generation date of the request. Statutory holidays were used in the list of holidays, which totalled to 10 days a year. Holiday dates were normalized by adjusting dates to observed holiday date. For example, if New Year's Day occurred on a Sunday, the observed date would be Monday, January 2.

Locations

Service request locations were published in two different formats – FSA code and street intersection. 15,847 unique street intersection descriptions were identified between January 2010 and June 2015. These locations were geocoded using a combination manual identification and the Google Maps Geocoding API, in order to identify their corresponding FSA.

A total of 14 street descriptions could not be geocoded, resulting in the exclusion of 500 service requests from the analysis. These location descriptions were too broad, making it impossible to determine which FSA it belonged to. For example, there are 121 service requests with the location description '*Don Valley Parkway N/Don Valley Parkway S|| Toronto*', which is a freeway which passes through 6 different FSAs.

District locations were also extracted based on service request location information. FSAs were mapped to their corresponding district, which was later merged with the service requests dataset. Requests using street names used the description after the '||' characters to extract the district information.

Weather Data and Events

Preliminary weather event analysis was done using Pearson Airport weather data and Etobicoke service requests as this was the only station that had detailed data outlining weather conditions, precipitation amounts, and temperatures.

Heat Warning events were grouped into Heat Warnings and Extreme Heat Warning categories, where an upgraded or downgraded warning was counted towards the level the warning was moved towards. For example, a downgraded Extreme Heat Warning was counted as a day with a Heat Warning. Cold Warnings are issued by the Toronto Medical Officer of Health when Environment Canada reports temperatures that are -15°C or colder, or -20°C or colder with wind chill[1].

Event data for cold warnings were extracted from the Pearson Airport hourly weather datasets based on these conditions. All other weather events were extracted from the weather description. A day was considered to have a weather event if it occurred at some point during the day.

For the correlation analysis, three days' worth of weather data was included for each service request - this includes the current day weather, as well as for the previous two days. This was done to account for any delay in response between a weather event, and when a service request is made. This was done by duplicating daily weather records, offsetting the date as needed, and merging everything together based on date. Total rain, total snow, and snow on ground values were unavailable for the Toronto City weather station, and was interpolated based on the average between the other three weather stations that were sourced.

Type Groupings

The 311 Service Request type dataset available on the Toronto Open Data Portal shows a total of 371 different classifications of request. However, the data set mentions that it is from November 2010 – which has since increased to over 500 unique instances of service request descriptions, as seen in the created service requests data set.

Service requests created after November 2010 were assigned a division, unit-section, and temporary code based on the formatting of the request description. There were also several request types that were duplicates of each other with minor variations to the description, such as usage of capitalization or spacing. These variations were corrected, and all versions were assigned the same attribute values.

In order to better manage the analysis, a taxonomy of service request types was created to group similar requests into categories. A total of 64 categories and subcategories were created, and the main 16 categories were used in the correlation analysis.

Feedback and publication type requests were analyzed without grouping counts into FSA since these types of requests do not have an address associated with them, and use '*King and John Street*' at the default location.

Table 1: The main categories that service request classes were grouped under, and the number of problem codes under each grouping.

Category	Description	Codes
Bylaw Enforcement	Enforce bylaws/regulations/licenses, such as noise complaints or property standards	16
City Infrastructure	City infrastructure and property that is not damaged, such as locate and relocate requests	83
Cleanup	Spill, debris, litter cleanup	13
Feedback	Comments/complaints about department staff, operations, and accessibility	44
Forestry Inspection	Tree and permit inspection, and removal operations	14
Infrastructure Damage	Reporting damaged city infrastructure and property	51
Infrastructure Maintenance	Request for maintenance or cleaning of city property and infrastructure, utility cuts, etc.	44
Miscellaneous	Miscellaneous requests	3
Publication	Requests for information in department publications	7
Registration	Opening an account with city services	2
Seasonal	Urban forestry related, such as tree planting and pruning	7
Testing	Water quality testing	3
Waste Collection	Missing or request for additional waste collection	120
Waste Receptacle	Damaged or missing waste receptacles	69
Waste	Waste request not related to collection or receptacles	2
Winter	Snow removal and salting/sanding due to ice	19

5. Analysis

In 2014, 311 Toronto received a total of 1.18 million calls, where 32.6% of them resulted in the creation of a service request. There has been an upward trend in 311 request volume every year since 2010 at an average rate of 13% a year.

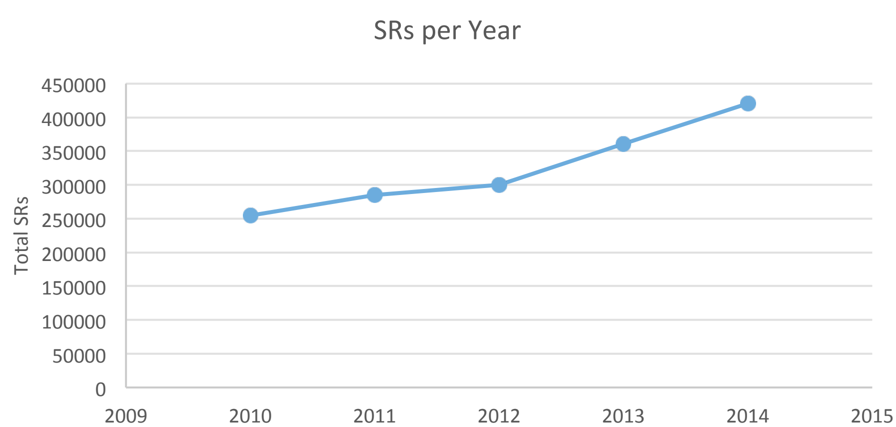

Figure 1: Yearly breakdown of received service requests.

What requests are being made?

From January 2010 to June 2015, *Solid Waste Management* made up the bulk of the requests averaging at 44.75% of all requests, while *Urban Forestry* makes up the least at 8.01%. However, as *Urban Forestry* primarily operates from the spring to the fall, *Municipal Licensing & Standards* makes up the next smallest division at 11.25% requests a year.

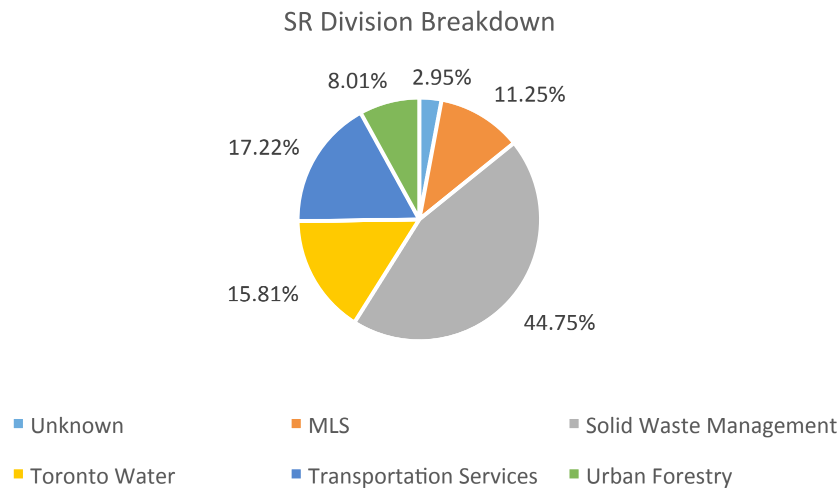

Figure 2: Breakdown of service requests from January 2010 to June 2015 grouped by the division that was responsible for the request. The 'Unknown' series are the service requests that could not be matched to a division based on the data provided.

Where are requests being made from?

Toronto can be broken down into 4 districts: *Old Toronto and York*, *Etobicoke, Scarborough*, and *North York*. For requests over the entire data set, when normalized by population, *Old Toronto and York* sees the highest request volumes at 0.81 per person, while *North York* sees the fewest at 0.58 per person.

Toronto can also be broken down into 102 unique FSAs. Overall, each FSA sees an average of 3058.5 service requests a year across all types. The FSAs with higher populations tended to have higher service request counts.

When are requests being made?

When looking at call volumes on a daily basis, there are two distinct volumes in service requests – one during weekdays, and one during weekends. Weekdays see an average of 1142.2 requests a day, while weekends see an average of 330.0 requests a day. Holidays see service request volumes at an average of 327.9 a day for both adjusted and non-adjusted dates, which is similar to weekends.

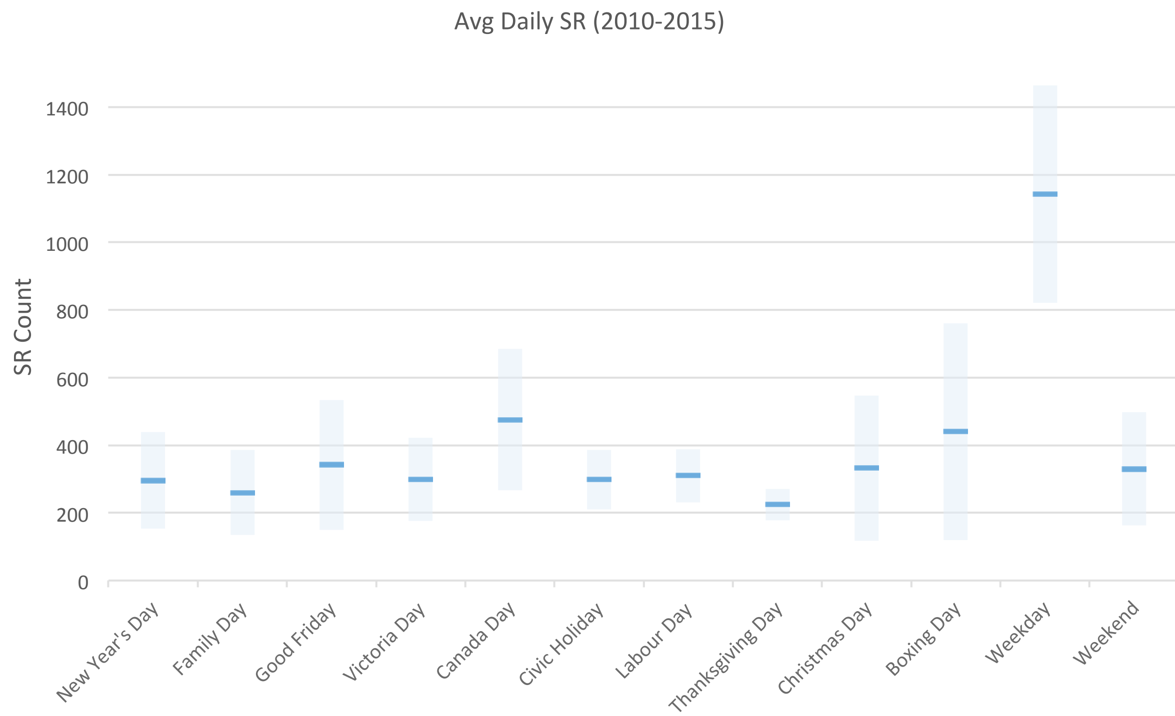

Figure 3: Average daily service request count grouped by type of day, where standard deviation is shown as the error bars.

During the course of a day, call volumes peak at 10-11 AM, with a sharp increase in requests 3-4 hours before hand. Weekdays see an average peak of 86.1 requests an hour, while weekends and holidays see a volume of 29.6 requests an hour. Weekday request volumes then drop to 76.6 an hour in the afternoon, and steadily decline to below 10 requests an hour between the hours of 3 and 11 PM. Weekends see a steady decline to below 10 requests an hour between 11 AM and 11 PM.

Over the course of the year, the summer months generally see the highest volume of service requests, and winters see the lowest.

How are calls being made?

While there has been significant growth in using online mediums since their inception in 2012, requests are still predominantly received through via phone calls. In 2010, 90.2% of all requests were from calls, which has dropped to 83.6% in 2014. Online submissions meanwhile only made up 1.2% of all requests in 2010, and has since grown to 8.4% in 2014.

Weather Correlations

The analysis started with a review of how weather events impacted service request volumes, followed by an in depth analysis comparing service request volumes to quantifiable weather descriptors, such as rainfall.

Weather Events

Average daily counts for clear and/or cloudy days was used as a baseline to represent days with no weather events. The average daily service request counts for each weather event type was then then compared against the baseline. While there was a 100-300 increase in daily service request volume on days with weather events, this was also accompanied with a standard deviation of 400 service requests. Weather event daily counts were also analyzed based on season. While this resulted in a more noticeable effect on daily averages for specific weather events (i.e. Heat Alert during the summer), this did not significantly reduce the variance in daily count.

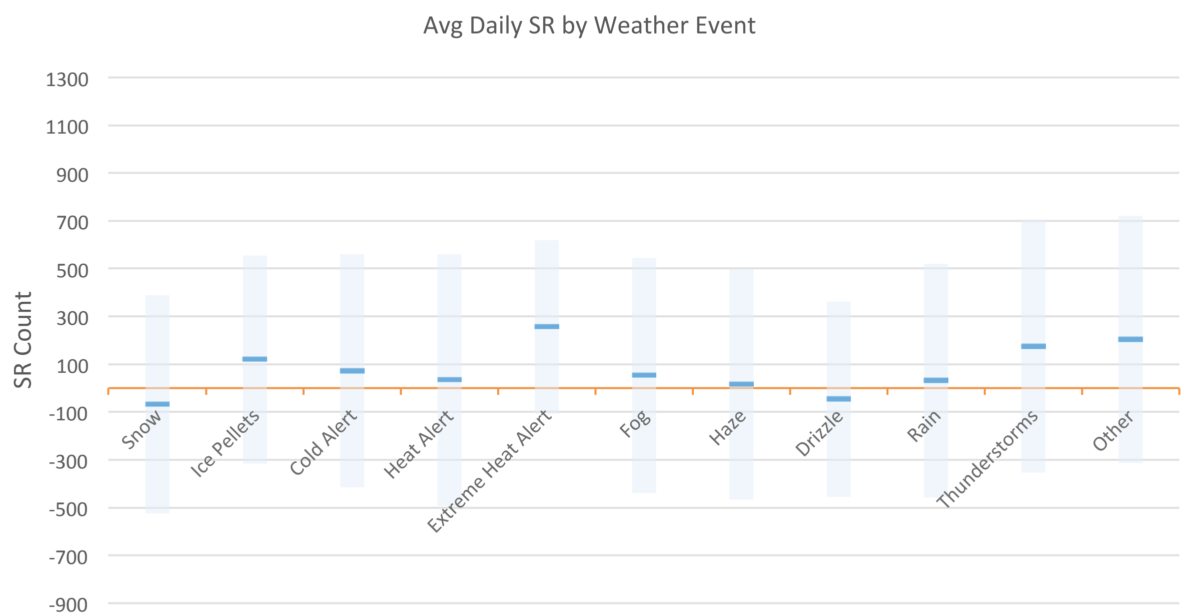

Figure 4: Average daily service request count grouped by weather event, where standard deviation is shown as the error bars.

Correlation Analysis

Pearson correlations were run between daily service request volumes for each FSA and its corresponding weather condition. The analysis compared daily service request volumes to mean daily temperature, total precipitation (in mm), total rain (in mm), total snow (in cm), and total snow on ground (in cm). Three variations of these comparisons were done against weather conditions on different days – on the day of the request occurred, as well as the two previous days. In addition, daily counts were split up based on request type. To visualize all the correlations, a heat-map was used, where each section was a different FSA. This resulted in 15 correlation heat-maps for each of the 17 request type groupings.

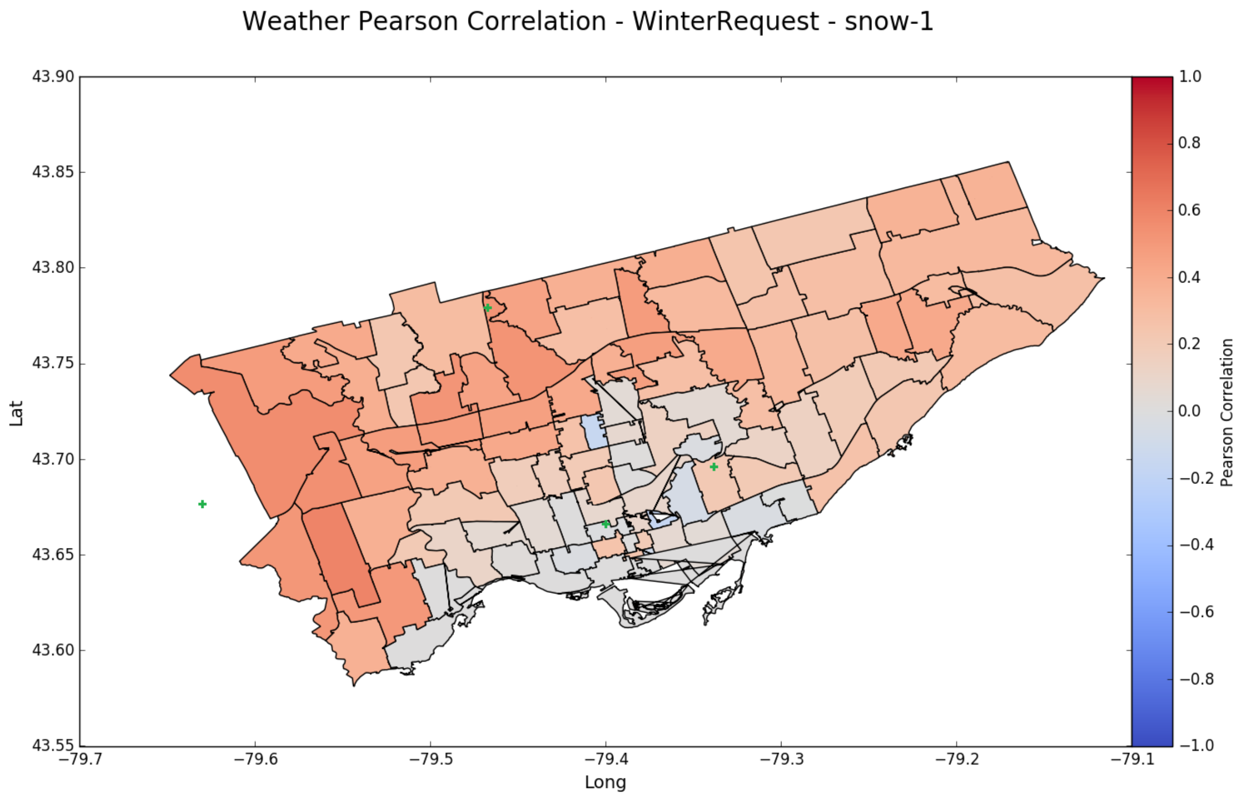

Figure 5: A heat-map showing the Pearson correlation between the daily service request count and snowfall from the previous day, where each section is a FSA. The four weather stations are shown as green crosses.

Most pairings shows weak correlations between 0.2 and 0.4. A handful of pairings, showed moderate correlations between 0.5 and 0.7. For example, in the heat-map above, the FSA ‘M9B’ shows the highest Pearson correlation value of 0.596 – one of the highest seen overall (see figure 6 below). . However, reviewing the data in closer detail confirmed there wasn’t a strong relationship between the variables Pairings with correlations between 0.6 and 0.7 had very few service requests, which skewed the correlation.

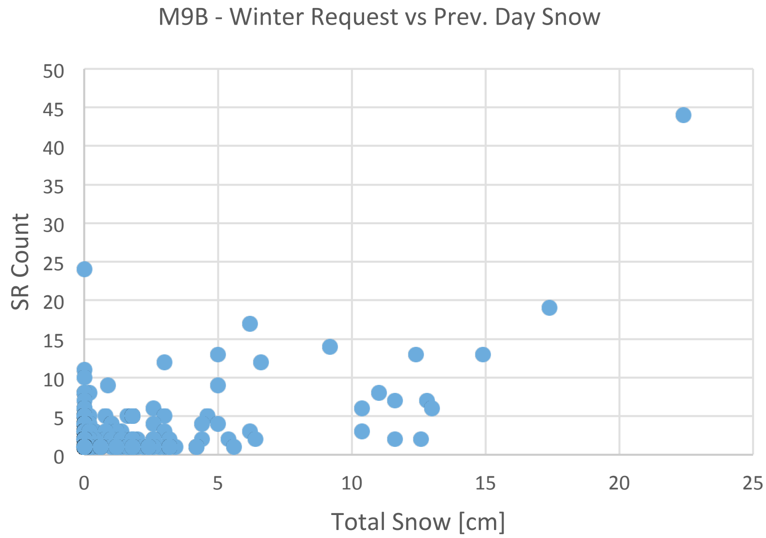

Figure 6: The scatter plot above shows the comparison between total snowfall and the daily count of all winter SRs at FSA 'M9B'. There are a total of 211 days which in which this FSA received winter SRs, and averaged 2.313 per day. The Pearson correlation was calculated to be 0.596, one the highest seen in amongst all weather-service request pairings.

Table 2: The daily average number and standard deviation of SR counts per FSA per category. Daily counts exclude days where the SR was not observed (i.e. Winter requests during summer months).

Category	**Average Daily SR Count by FSA**	**Average Daily Std Dev by FSA**
Bylaw Enforcement	2.04	1.77
City Infrastructure	1.81	1.26
Cleanup	1.17	0.44
Feedback	1.43	0.78
Forestry Inspection	1.68	1.80
Infrastructure Damage	1.77	1.19
Infrastructure Maintenance	1.67	1.54
Miscellaneous	1.00	0.00
Publication	1.34	0.80
Registration	1.25	0.81
Seasonal	1.35	0.68
Testing	1.14	0.42
Waste Collection	3.49	3.29
Waste Receptacle	2.98	2.38
Waste	1.00	0.00
Winter	2.64	3.04

6. Discussion

The analysis suggests that there is a weak direct relationship between weather and service request volume aside from the obvious cases, such as snowfall and salting requests. Even then, the strength of the relationships are moderate at best with only a few pairings showing correlations between 0.5 and 0.60. While the 16 categories used in the analysis were further broken down into additional sub-categories, analyzing it shows significantly weaker correlations due to insufficient amounts of data.

Counts were grouped by FSA in order to account for differences in geography, infrastructure, and management of operations. However, it is possible that different groupings of requests would yield better results. It is also possible that weather has an influence on another variable altogether which in turn has a direct effect on service request volume.

Inconsistent/Incomplete Data

There were several limitations with the data sets which made it difficult to perform analysis.

First, the location data for the service requests were published in two different formats. While it is understandable that addresses are anonymized to protect the identity of the customer, it would be significantly easier to explore if the format was consistent. While having all locations geo-coordinates would have been the most ideal, identifying the other geographical areas that the service request belonged to, such as ward or neighbourhood, would be sufficient to facilitate other angles for analysis. For example, datasets from the Wellbeing Toronto[6] initiative, such as demographics or economic indicators, are grouped based on neighbourhood.

There were also cases where datasets were not complete or up to date. The 311 Service Request type dataset is missing around a quarter of the request codes used in the 311 Service Request datasets. The inconsistent formatting between request description and codes between departments made it difficult to automatically extract common features.

7. Conclusion and Future Work

This paper investigated the relationship between the volume of 311 Toronto service requests and weather phenomena, in order to determine if predictive modeling techniques could be applied in order to optimize resources and improve customer service. Service requests were aggregated on a daily basis, grouped by FSA, and compared against the weather. Daily counts were compared against the occurrence of various weather events, as well as weather attributes, such as rainfall, snow, and temperature. Relationships between service request volumes and weather were found to be weak to moderate.

For future work, there are two main avenues to explore in terms of service request prediction – creating more customized models, and incorporating data related to business practices.

[6] http://map.toronto.ca/wellbeing/

The variety of service requests would likely result in the development of customized models for each type by incorporating additional data to explore different angles of analysis. For example, exploring the long term effects of winter weather and the damage it causes to infrastructure, in order to predict maintenance related request volumes. With this approach, data pertaining to road construction and maintenance, traffic patterns, and weather could be used to build a model to determine the frequency and severity of potholes, which can then be linked to the volume of pothole related service requests.

Another option to consider is that it is also possible there are moderator variables which the weather influences, which in turn, affects request volumes. For example, snow removal requests can only occur after or during a snowfall. However, rather than directly influencing service request volume, the amount of snow may influence scheduling, which in turn may correlate with service request volume. Incorporating business operation key performance indicators and schedules may lead to a better model of how weather influences service request volume.

Acknowledgements

We would like to acknowledge the support we received from Toronto 311 staff, including Gary Yorke and Abdi Osman.

References

Andrews, B. H., & Cunningham, S. M. (1995). L. L. Bean Improves Call-Center Forecasting. *Interfaces, 25*(6).

Avramidis, A. N., Deslauriers, A., & L'Ecuyer, P. (2004). Modeling Daily Arrivals to a Telephone Call Center. *Management Science, 50*(7), 896-908. doi:10.1287/mnsc.1040.0236

Brown, L., Gans, N., Mandelbaum, A., Shen, H., Zeltyn, S., & Zhao, L. (2005). Statistical Analysis of a Telephone Call Center: A Queueing-Science Perspective. *Journal of the American Statistical Association, 100*(469), 36-50. doi:10.1198/016214504000001808

Ibrahim, R., & L'Ecuyer, P. (2012). Forecasting Call Center Arrivals: Fixed-Effects, Mixed-Effects, and Bivariate Models. *Manufacturing & Service Operations Management, 15*(1), 72-85. doi:10.1287/msom.1120.0405

Jongbloed, G., & Koole, G. (2001). Managing uncertainty in call centres using Poisson mixtures. *Applied Stochastic Models in Business and Industry, 17*(4), 307-318. doi:10.1002/asmb.444

Kungle, R., & Maluchnik, J. (1998). Call Center Forecastingat AAA Michigan. *Journal of Business Forecasting, 16*(4), 8-13.

Millán–Ruiz, D., & Hidalgo, J. (2013, November). Forecasting Call Centre Arrivals. *Journal of Forecasting, 32*(7), 628-638. doi:10.1002/for.2258

Pacheco, J., Millán–Ruiz, D., & Vélez, J. L. (2009). Neural Networks for Forecasting in a Multi-skill Call Center. (D. Palmer-Brown, C. Draganova, E. Pimenidis, & H. Mouratidis, Eds.) *Engineering Applications of Neural Networks, 43*, 291-300. doi:10.1007/978-3-642-03969-0_27

Schwester, R. W., Carrizales, T., & Holzer, M. (2009, July 01). An Examination of the Municipal 311 System. *International Journal of Organization Theory and Behacior, 12*(2), 218-236.

Shen, H., & Huang, J. Z. (2008). Interday Forecasting and Intraday Updating of Call Center Arrivals. *Manufacturing & Service Operations Management, 10*(3), 391-410. doi:10.1287/msom.1070.0179

Taylor, J. W. (2008, February). A Comparisonof Univariate Time Series Methods for Forecasting Intraday Arrivals at a Call Center. *Management Science, 54*(2), 253-265.

Weinberg, J., Brown, L. D., & Stroud, J. R. (2007). Bayesian Forecasting of an Inhomogeneous Poisson Process with Applications to Call Center Data. *Journal of the American Statistical Association, 102*(480), 1185-1198. doi:10.1198/016214506000001455

Zha, Y., & Veloso, M. (2014, June). Profiling and Prediction of Non-Emergency Calls in New York City. *Workshops at the Twenty-Eighth AAAI Conference on Artificial Intelligence*.

