Maximizing Coverage of EMS services in New York City using a Double Standard Model

Henry Dinhofer
Courant Institute of Mathematical Sciences, New York University
New York, NY
hld259@nyu.edu

Prateek Sappadla
Courant Institute of Mathematical Sciences, New York University
New York, NY
pvs266@nyu.edu

Lakshminarayanan Subramanian
Courant Institute of Mathematical Sciences, New York University
New York, NY
lakshmi@cs.nyu.edu

ABSTRACT

We incorporate a deterministic modeling algorithm, the Double-Standard Model (DSM) to recommend optimal ambulance placement over New York City. Then we tested our model on several different days in 2016, with data from NYC EMS Ambulance dispatch. The DSM model is able to place ambulances in all locations and come up with ambulance placements that satisfy 95% coverage or better on both the lowest demand day and highest demand days of 2016. We modified DSM to include the capacity constraint which accounts for the number of cases an ambulance can serve in a particular time frame to better encapsulate demand for an area. We build flexible models that can be more specifically customized to the various needs of ambulance placement.

KEYWORDS

Resource Allocation, Ambulance Placement, Double Standard Model, New York City, Maximal Covering Problem

ACM Reference Format:
Henry Dinhofer, Prateek Sappadla, and Lakshminarayanan Subramanian. 2018. Maximizing Coverage of EMS services in New York City using a Double Standard Model. In *Proceedings of ACM SIGKDD conference (MUD3).* ACM, New York, NY, USA, Article 4, 7 pages. https://doi.org/10.475/123_4

1 INTRODUCTION

Ambulance response times matter. Placing emergency medical ambulances in the right spot is critical to reducing these times. We built an ambulance placement machine that places ambulances, in the context of New York City. Our machine utilizes a modified version of the Double Standard Model (DSM) proposed by Gendreau et al. This allows for a flexible, scalable, unbiased and more reliable model.

In New York City over 7,000,000 ambulance drives have been dispatched over the past 5 years. Recent research in medical journals indicates that the time these ambulance rides take to reach a patient has been directly correlated to health outcomes[9, 13]. We decided to test out placement of ambulances on a map, utilizing the Double-Standard Model (DSM), a model that incorporates standards of coverage as well as cooperation, capacity and specificity with sub-coverage elements.

The model is flexible because it can work with varying numbers of ambulances. It is scalable because works with hundreds of demand points. It is unbiased because democratizes resources to focus on all demand and not the few high demand points. It allows a third party to set the criteria for a given day, and computes the entire ambulance placement for the city.

We've introduced a capacity constraint k that certifies demand can be achieved in a region. This certification of a region is amongst several adjacent regions a local maxima, that more comprehensively addresses demand factoring in nearby ambulances.

The need for decreased ambulance response times coincides with the need for optimal ambulance placement. Ambulances if placed in the right position, will be able to cover more cases while also being able to reach them faster and within the current legal standards. The system we built allows for wider levels of flexibility, is just as durable as existing solutions and could be more fine tuned to meet the various needs and demands of the Fire Department of New York's (FDNY) Emergency Medical Services (EMS) system.

From this study we discovered important trends in NYC data such as seasonality, weekday, time-of-day and can use these insights into applying our modeling to calculate placements for any given day. To the best of our knowledge, this is the first study that has looked at New York City EMS data this way.

2 RELATED WORK

2.1 Problem

Optimal ambulance placement has been studied for several years, and various approaches have evolved. Eaton et al. first prescribed ambulance placement; using a Maximal Covering, they calculated placement using a form of the Maximal Covering Location problem (MCLP). [6]

In the years since Eaton, optimal ambulance placement has undergone three generations of thought: (a) Deterministic, (b) Probabilistic and (c) Dynamic modeling.

Deterministic thought is any ambulance modeling that assumes coverage is reachable within a circular zone of driving time, and once an ambulance's location is optimally placed, the vehicle is set in a static immovable position for the rest of the day. P-centering, Maximal Covering Location Problem (MCLP), P-median solutions all contributed to tenable applications for optimal ambulance placement [4, 10]. This coverage thinking has been continuously improved upon. Recently the Double Standard Model (DSM) was proposed by Gendreau et al.[7] and has since been implemented in regions in Belgium and nearly all of Austria, [12] as well as the cities of Vienna[14], Montreal [12] and Tijuana [5]. Subsequent analysis of these coverings has also taken call data and looked into whether time-of-day trends affect ambulance placement.[14, 15] We agree that patients move their location during the day and as such our ambulances need to move accordingly to respond to that shifting demand. While initial results have been inconclusive, given more accurate call locations we think that specific sub-regionality trends can emerge. Dibene et al.'s [5] paper aggregates time-of-day over a yearly window. In our initial analysis we found out that both monthly seasonality and also weekly trends such as weekday or weekend call volume did appear.

Probabilistic thought groups ambulance placement around the use of a percent likelihood that represents: an ambulance's likelihood it is unavailable because it has already been assigned to a call, or a percent likelihood a demand can be reached within time t, or factoring in neighborly cooperation as a percentage of signal strength such as described in Berman et al.'s *Cooperative Maximum Covering Location Problem* (CMCLP) [3, 11]. Berman work has not yet been applied to the ambulance placement problem specifically but his maximal covering for the placement of sirens on a map is very similar to the optimal ambulance location application. There is an opportunity here to explore more into overlapping and cooperative set coverings. Both ambulance unavailability and cooperation were the inspiration for our capacity *(k)* and cooperative thinking see *Section 4*.

Dynamic shifting is the newest trend in ambulance modeling. This is applied to realtime and live reorganization. When an ambulance is no longer available to cover an area due to assignment, additional processing is done to determine the value of "shifting" the other waiting ambulances to better respond to demand. Using estimates of future demand, the above mentioned probabilistic modeling of ambulance busyness is incorporated into determining whether an ambulance should be moved. Dynamic ambulance shifting has been done with a modified DSM to make a *Dynamic Double Standard Model* (DDSM) [8]. Additionally studies coming from the Netherlands point to decreases in response time on average of 16-20%.[16] While these are promising and encouraging results, we didn't elect to dive into dynamic ambulance placement as the current NYC system uses a deterministic and not dynamic system for its Basic Life Support (BLS) units, the significant additional time to model such a system, as well as have higher costs than the existing system. This is based off of intuition and conversations with current FDNY personnel. Knowing that the FDNY, a public government department, is reluctant to increase costs we elected to approach the problem with an "as is" approach. Assessing the higher expenses from higher gas usage and the higher need for vehicular repairs, is beyond our current knowledge of the FDNY EMS System.

2.2 Our baseline - Dibene et al [5]

Our study is an application of the DSM model to New York City data for two different days. We looked to successful applications in the cities of Vienna Austria, Montreal Canada, and Tiujana Mexico. In particular Dibene et al.'s[5] work in DSM incorporated analysis among four different time-of-day windows. They aggregated all city-wide calls in 2014 to these time windows. Furthermore they also made distinctions between working week (weekday) and off-day (weekend), taking the four windows above and applying them again making in total eight different time scenarios. Most calls happen on weekdays during the afternoon and evenings.

The Dibene et al. paper also used a relatively low number of calls (n = 7,746) for the entire year. This is very similar to the total call volume in just our two days (high demand + low demand = 7,500). Using the DSM modeling they calculated placements for number of ambulances p (p = 6 to 22). Of note they calculated coverage as existing within the set covering circle. Dibene's team however did not note whether this set covering would be adequate enough for any given individual day. We wanted to test our covering on two individual days, and not on an aggregate basis. Before an actual application of DSM to New York City ambulances, we would test out the various levels of efficacy over aggregated daily, weekly, time of day, monthly and annual aggregate demand.

3 EMS DATASET

We applied DSM to the EMS Incident Dispatch Dataset [2], selecting two different days to simulate the model on. One was on a day of high demand where 4700 calls came in, and the other was on a low demand day with 2800 calls. The NYC OpenData set consists of 7,023,225 ambulance assignments, over a period of 5 years from January 1st 2013 to December 31st 2017. The term "dispatch" means any time a 9-1-1 operator during the course of a call, makes the decision to send an ambulance to drive to a demand point. All demands were aggregated on the zipcode level, there were concerns about potentially exposing patient information that OpenData was circumspect with possibly breaching HIPAA compliance. All W ambulance sites were aggregated to the zipcode level.

Analysis of the EMS Incident Dispatch Dataset was performed using pandas an open source Python package for large data analysis, and Jupyter notebook for compilation. Data for the two days was then fed into our model that processed the coverage using the python linear programming tool PuLP. It returned a list of zipcodes and the number of ambulances at that zipcode that ambulances would be placed at to provide optimal coverage. It also returned the percent of coverage the placement achieved.

The Arc GIS (Geographic Information System) mapping tool was used to map the frequency counts of calls by zipcode to a map of New York City. Zipcode polylines for Arc GIS were referenced by its earliest known locations from 2010 city and Baruch College data. To calculate the travel time between two zipcodes required converting zipcodes to GPS coordinates of the centroid of the zip. GPS Location points for the centroid of the zipcode were referenced from 2013 Government Data. [1]. The GPS coordinates of these centroids were then used to calculate travel times t_{ij} between two zipcodes via the *Google Distance Matrix* API. While the Distance Matrix tool provided the ability to forecast travel times at specific times in the day, the authors of this paper decided to go with calculating travel times at 2:00pm on a Monday as additional overhead was required, and the predictions tool was prone to making subtle mistakes such as giving a nonzero number for driving time as the time between

Variable	Description
V	Set of demand points. Zipcodes in our model
W	Set of ambulance sites. Zipcodes in out model
t_{ij}	Travel time between point i and point j
x_j	Number of ambulances placed at demand point j
r_1	First time standard
r_2	Second time standard
p	Total number of available ambulances
k	Number of calls serviced by an ambulance in a day
$W_i^{r_1}$	$\{j \mid t_{ij} \le r_1\}$
$W_i^{r_2}$	$\{j \mid t_{ij} \le r_2\}$
d_i	Number of calls(demand) at demand point i
z_i	$\sum_{j \in W_i^{r_1}} x_j$
y_i	$\sum_{j \in W_i^{r_2}} x_j$

Table 1: Description of Variables used in the model

demand point i and ambulance site i, to the same point $t_{ii} = 0$ cannot have any travel time.

4 MODEL

We use a variation of the Double Standard Model[7] to determine an optimal placement for the ambulances in New York City. The terminology used to define the model is described in TABLE 1.

The problem we are looking to solve is that of placing p ambulances optimally across the available ambulance sites W such that we obtain maximal coverage within a time standard r_1 of the demand across the set of demand points V. Additionally, all demand should be adequately covered within a second more lenient time standard r_2.

The capacity available to serve a demand point i is given by the product of the number of ambulances within a time radius of r_1 and the average number of cases that can be served by an ambulance which is given by k. We consider a case to be adequately covered if and only if there is a capacity of 2 available to serve it. Generally, an ambulance is within r_1 time range of multiple demand points. Thus it is considered as being available to serve the calls originating out of multiple demand points. The double coverage constraint is placed to prevent the model from overestimating the capacity of the ambulance placement that may arise due to the same ambulance being available for multiple demand points.

The introduction of the capacity constraint k is an important addition of our work over the standard DSM model. The standard DSM model considers all cases at a demand point to be covered if there is an ambulance within time r_1 from the demand site. We feel this does not sufficiently capture the complexity of the real world scenario. This assumption will break in the case when there are much more calls originating from a demand point than what a single ambulance can serve. The capacity constraint k provides a simple but better way of modeling the capacity available.

The variable z_i represents the number of ambulances available to serve the demand point i within a time period r_1. Similarly, y_i denotes the number of ambulances available to serve the demand point i within a time period r_2.

We are looking to maximize the coverage that can be covered atleast twice within a time standard of r_1 while ensuring that atleast $\alpha\%$ of the demand in each zipcode is covered twice within r_1 and all demand is covered atleast once within a second time standard r_2. The optimization problem for maximizing the coverage by optimal placement of ambulances can then be formulated as below:

$$Maximize\ \Big(\sum_{i \in V} min\left(d_i\ ,\ \frac{kz_i}{2}\right)\Big)$$

subject to the constraints

$$\forall\ i\ \in V,\ d_i\ \le ky_i \tag{1}$$

$$\forall\ i\ \in V,\ \frac{kz_i}{2}\ \ge \alpha d_i \tag{2}$$

$$\sum_{j\ \in\ W} x_j\ =\ p \tag{3}$$

$$0\ \le x_j\ \le p_j \tag{4}$$

$$\forall\ j\ \in W,\ x_j\ is\ an\ integer \tag{5}$$

The objective function is a summation over all demand points V. For each demand point $i \in V$, we take the minimum of the demand at that point (d_i) and the coverage capacity available to serve that point. This minimum is to ensure that if at a particular point, the available capacity is more than the demand at that point, we do not over score our model for that point. The constraint (1) is the second covering constraint given by the double standard model which enforces that all demand is covered once within a second time standard r_2. The constraint in equation (2) ensures that atleast $\alpha\%$ of the demand at each demand point is covered by the placement of the ambulances. This constraint prevents the concentration of ambulances in a few places with very high demand. Constraints (3) and (4) ensure that a total of p ambulances are used and at most p_j ambulances are located at site j respectively.

5 RESULTS

This section describes the coverage values obtained for different values of the parameter values α and p_j for two different days representing high and low demand days. Table 2 details the coverage values for a high demand day with $\alpha = 0.95$ and $p_j = 5$. Table 3 documents the coverage obtained by increasing p_j to 10 while keeping α the same. It shows an interesting scenario whereby increasing p_j allows us to satisfy demand at a lower cumulative capacity. The coverage values for $\alpha = 0.7$ for a high demand day are detailed in Table 4. Finally, Table 5 shows that the demand can be satisfied with a much smaller number of ambulances on a day of low demand. The above mentioned results are described in further detail in the following subsections.

5.1 DSM

The United States EMS Act prescribes that 95% of the demand should be covered within 10 minutes. With this in mind, we set $r_1 = 10mins$ and $\alpha = 0.95$. The second time standard within which all demand is to be covered was set to $r_2 = 15mins$. The maximum number of ambulances at each site was restricted to $p_j = 5$. Lets call this initial configuration of the parameters as $C1$. We tried different combinations of values for the total number of ambulances p and the daily capacity of an ambulance k. The total number of calls can vary greatly depending on the day. On certain days, when the demand is highest, the total number of calls in a day was observed to be as high as 4700. On a day when the demand is lowest, the number of calls was observed to be around 2900. The results for the configuration $C1$ of the parameters for a day with the highest demand are given in Table 2.

In Table 2, we observe that the cumulative capacity of the ambulance system for values $p = 250$ and $k = 21$ is less than that for the values $p = 400$ and $k = 15$, a $250 \times 21 < 400 \times 15$. However we observe that although the problem is feasible for the former configuration, we get an Infeasible result for the latter. This behavior can be understood by looking at Figure 1.

Figure 1 is a histogram with the X-axis representing the number of calls in a day and the Y-axis representing the number of zipcodes. We can see from the plot, that there are a few zipcodes with very high call volume whereas most of the zipcodes have very few calls. The capacity at each ambulance site is restricted by the maximum number of ambulances available at the site. In configuration $C1$, this is restricted to 5. Thus for a lower value of k, the capacity at any of the sites is not enough to meet the constraint (3) for the high demand zipcodes which results in an Infeasible state. This can be remedied by increasing the maximum number of ambulances allowed at each site. Table 3 shows the results obtained by setting $p_j = 10$.

In order to test the flexibility of our model, we tried a configuration $C2$ with $\alpha = 0.7$. Table 4 details the results for the configuration $C2$. By decreasing α we are able to cover more than 92% demand with just 250 ambulances. However, by comparing Table 3 and Table 4 it can be seen that for higher number of ambulances and higher value of k, setting a lower value of α degrades the performance. This degradation in performance could be attributed to the fact that with a lower α, the model becomes more lenient and ends up placing more ambulances in zipcodes that are more spaced out with fewer overlapping coverages. Thus making the optimization problem less strict actually degrades the total coverage obtained.

Table 5 describes the results of running the model with configuration $C1$ on a day with the lowest observed number of calls. It can be seen that on days with low demand, a very high coverage can be obtained with much fewer ambulances. Thus based on the demand expected for a day, the total number of ambulances deployed in the city can also be optimized using this model.

5.2 Comparison with Dibene et. al

The Double Standard Model as applied in Dibene et. al is our baseline comparison. As mentioned before, they segmented ambulance placements by eight different times segments during a week. Their formulation remains consistent with prior DSM formulations, the goal being to maximize the number of demand points reachable at least twice within a given time standard r_1. This r1 is represented as a yes or no action. Either a demand point exists within a circular range of coverage or it does not. This leads to inflated expectations about an ambulance's ability to respond to a call. For a given day, our DSM shows the actual coverage available, not a loose grouping of covering whereby we assume an ambulance or two can achieve an extraordinary feat such as meeting more than > 70 calls. We propose a realistic covering.

K provides an upper bound to the demand that is realistically reachable. Either it is just the demand that is serviceable or the total demand at demand point d_i. Our model differs from Dibene's in that it uses the raw number of EMS calls to calculate coverage, and not a weighted sum. Thus we can count an accurate measure of EMS calls that are reachable.

5.3 How many ambulances do you need?

The p number of ambulances is initially set at the start of our model. From our results, p increased and decreased depending on the overall demand that day, as well as the capacity k for each ambulance.

5.4 What is the impact of changing k?

Changing k positively increases potential coverage. The more demands that an ambulance can reach the more total coverage it will provide.

5.5 What is the impact of changing p?

In general, increases in the number of ambulances provided increases coverage. In Table 3 there isn't an increase in coverage due to the value of p_j constraint (4). Ambulances can be placed at sites but they are limited by the maximum number placeable for any given site. This tells us that $p_j = 10$ ambulances wasn't enough to meet demand at a high demand point.

5.6 Additional Results

Initial analysis revealed a few trends in New York City's emergency medial service. As described by Dibene et al., time of day directly affected the number of calls that are serviced. The authors of this paper experienced similar ratios for the number of calls during the four time sub- segments (Figure 2). What was most prevalent was that the calls occurred most in the afternoon, followed by evening, then morning and finally as EMS professionals call it, the "Third Watch" that is the watch that happens late at night.

Differences between the day in the week were minimal. Figure 3 describes each day of the week and the number of calls for each day. There is a 7% difference in weekday calls and off-days /weekend calls, Δ n = 73,775. This is a minor factor in ambulance demand.

We also looked into possible seasonality trends. Figure 4 describes the breakdown of calls over a monthly window, and while there is a noticed difference between mid-year (most) and winter months (least), we could only establish broader seasonality trends. From the highest month (July) to the lowest (February), we found a 16% difference, Δ n = 96,156.

Finally geographic analysis with Arc GIS (Figure 5) shows that there is significant clustering of calls in a few regions. This is weighed towards a handful of zipcodes out of the 209 possible zipcodes. Since the frequency with which a few zipcodes call far outweighs that of most others, it is clear that demand is highly concentrated amongst a few zipcodes.

		Number of Ambulances (p)				
		250	*300*	*350*	*400*	*450*
Daily Capacity of Ambulance (k)	*15*	Infeasible	Infeasible	Infeasible	Infeasible	Infeasible
	18	Infeasible	Infeasible	Infeasible	Infeasible	Infeasible
	21	99.80716	99.81787	99.81787	99.86072	99.86072
	24	99.89286	99.89286	99.89286	**99.91429**	**99.91429**

Table 2: Percentage coverage with parameters settings as follows. $r_1 = 10mins$, $r_2 = 15mins$, $\alpha = 0.95$, $p_j = 5$. **High Demand day**

		Number of Ambulances (p)				
		250	*300*	*350*	*400*	*450*
Daily Capacity of Ambulance (k)	*15*	Infeasible	Infeasible	99.70002	99.7429	99.7750
	18	99.7643	99.7643	99.7643	99.7643	99.7643
	21	99.80716	99.81787	99.81787	99.81787	99.81787
	24	99.89286	99.89286	99.89286	**99.91429**	**99.91429**

Table 3: Percentage coverage with parameters settings as follows. $r_1 = 10mins$, $r_2 = 15mins$, $\alpha = 0.95$, $p_j = 10$. **High Demand Day**

		Number of Ambulances (p)				
		250	*300*	*350*	*400*	*450*
Daily Capacity of Ambulance (k)	*15*	90.36855	90.36855	90.61496	91.53632	92.4683
	18	90.8078	90.8078	91.23624	92.73623	93.0148
	21	92.2327	92.43625	92.69338	93.7433	93.8397
	24	92.37197	92.82194	93.7433	94.40754	**94.7075**

Table 4: Percentage coverage with parameters settings as follows. $r_1 = 10mins$, $r_2 = 15mins$, $\alpha = 0.7$, $p_j = 10$ **High Demand Day**

		Number of Ambulances (p)	
		200	*250*
Daily Capacity of Ambulance (k)	*15*	99.82251	99.82251
	18	99.87575	99.87575
	21	100.0	100.0
	24	100.0	100.0

Table 5: Percentage coverage with parameters settings as follows. $r_1 = 10mins$, $r_2 = 15mins$, $\alpha = 0.95$, $p_j = 10$ **Low Demand Day**

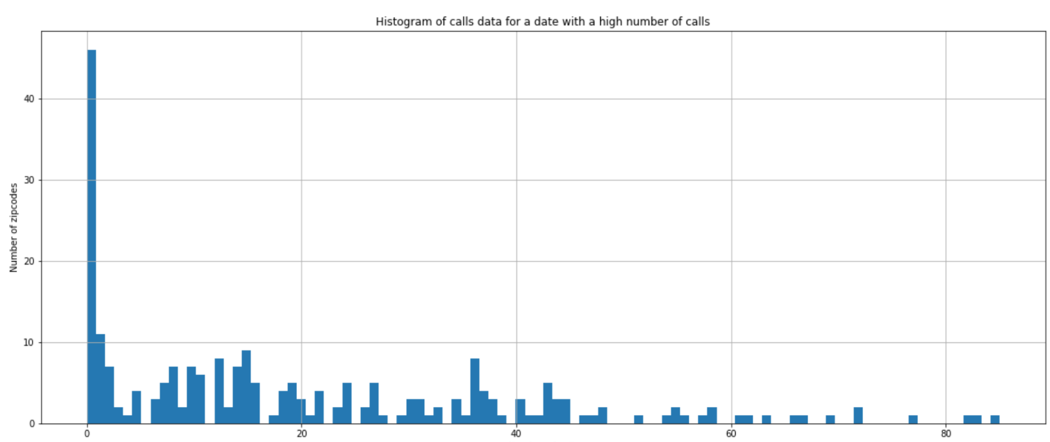

Figure 1: Histogram of count of calls on a high demand day

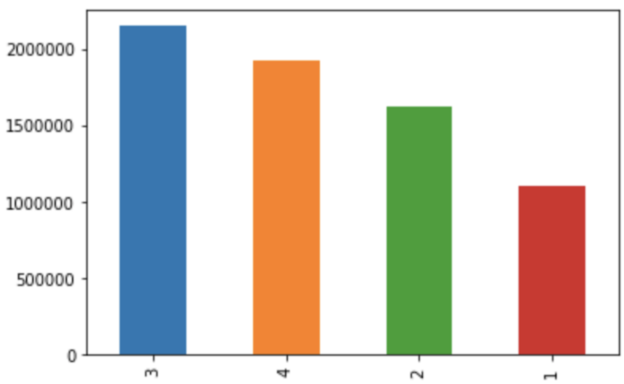

Figure 2: Time of day and count of calls during a 6 hour time interval.
3 = 12:00pm - 5:59pm (Afternoon), 4 = 6:00pm - 11:59pm (Evening),
2 = 6:00am - 11:59am (Morning), 1 = 12:00am - 5:59 am (Night)

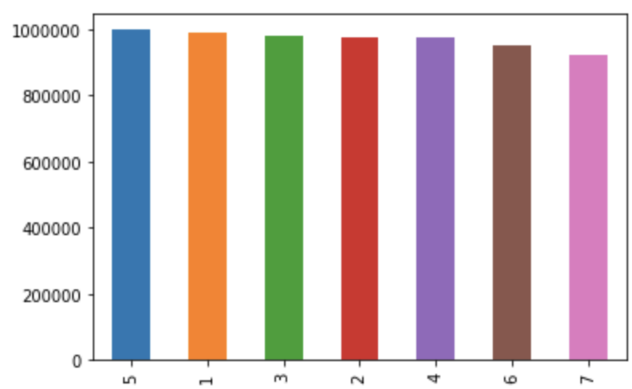

Figure 3: Day of the week and count of calls for that day
1 = Monday, 7 = Sunday

6 DISCUSSION

The DSM model allows for any given demand point to be covered minimally within time standard r_2, while *maximizing the demand covered at least twice* within time standard r_1. We were most surprised by the fact that on a low demand day, 100% of demand points could be reached by the double standard with as few as 200 ambulances. This contrasts greatly with demand at a high demand day, where more than 450 ambulances are required to meet existing U.S. EMS standards.

Seasonality, as represented by month-to-month needs could be factored into the broader number of ambulances needed for any given month.

While we tested our data on two pre-selected days, we would like to extend our model to work over aggregated data. As mentioned earlier this model could be used to test out the various levels of efficacy over aggregated daily, weekly, time of day, monthly and annual aggregate demand. It wouldn't take much to convert our model to a broader picture. We would have to multiply capacity k by some n number of days, but the remaining constraints would remain all within the modeler's control and discretion to fine tune.

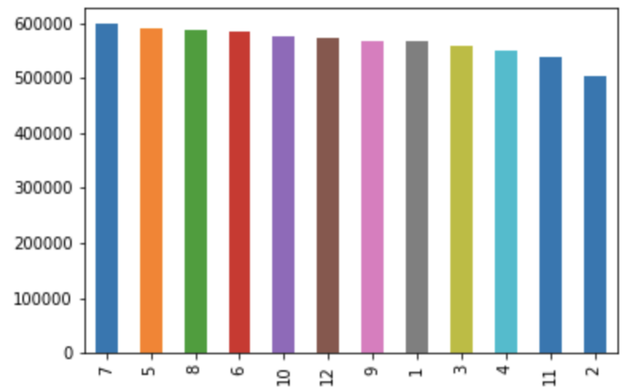

Figure 4: Seasonality trends, Counts of calls for a given month.
1 = January, 12 = December

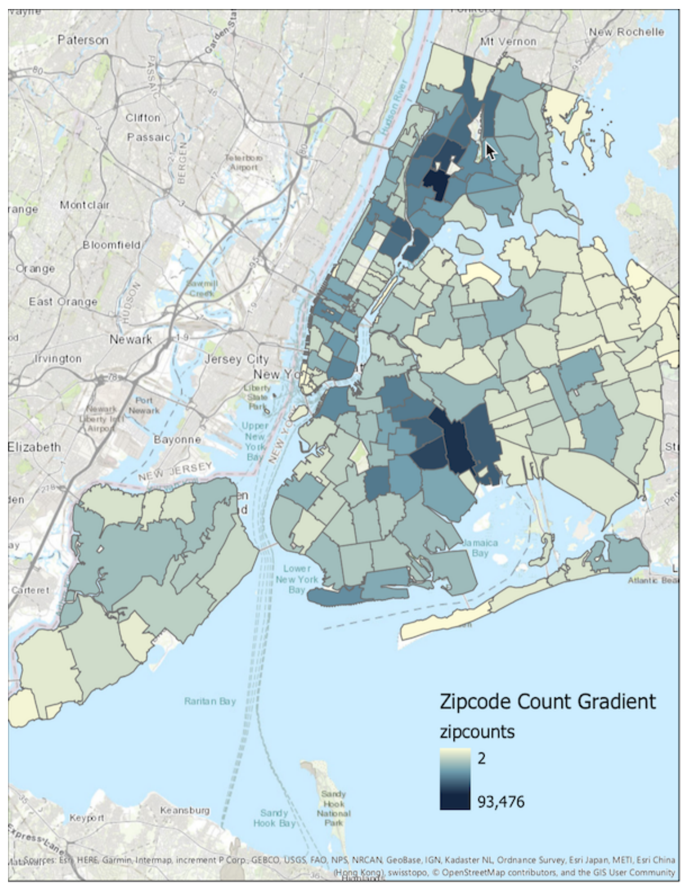

Figure 5: Number of call counts for every zipcode

Our study has several limitations. The location data of the ambulance calls available to us is only at the zipcode level. We do not have location data at a finer level. Zipcodes are too big a unit to get an effective placement of ambulances. It would be interesting to get location data of the demand points and a larger set of possible ambulance sites and then fit the same model for this finer level data. Also, the use of finer data points will also enable us to get better estimates of the travel times thereby leading to more accurate models.

We were also limited by the lack of knowledge of the current NYC EMS system. Current ambulance placement sites W were unknown so we were unable to establish a baseline to compare our results to.

7 CONCLUSIONS AND FUTURE WORK

We have modeled NYC EMS data using the DSM model for optimal ambulance placement. We tested it on two different days, and recommended two unique ambulance placement regimens for both of those days.

Another interesting extension of this model would be for a tiered ambulance system consisting of both Basic Life Support and Advanced Life Support Ambulances. The Basic Life Support ambulances would be required to cover all the calls while the Advanced Life Support ambulances would cater to high severity calls. A future task is to develop an extension of the double standard model with these two types of ambulances and different time standards for each type of ambulance.

In the future, we aim to simulate our results and compare them to the existing system, but as mentioned at the end of the last section we would need both more specific data about current ambulance placements, as well as specific location data (even if its generalized to a city block) for calls. This would then be run on a scheduling simulation. From there we could get specific numbers to whether response times are lowered, coverage is increased, and whether our model is definitively an improvement over the existing EMS ambulance placements. Eaton once said "simulations are great but... such models are descriptive rather than prescriptive." We agree with his sentiments as well. We would also like to possibly extend this to test out a Dynamic modeling. The most natural extension to this would be the Dynamic DSM (DDSM) mentioned earlier.

We also aim to run a prediction model for predicting future demand. This could incorporate such elements as seasonality. We initially attempted to run an LSTM to gather future predictions but it required more analysis and instead chose other avenues. It remains to be seen whether sub-regional seasonal trends may exist. The more accurate the predictions, the better the modeling, and the better the utilization of resources.

REFERENCES

[1] 2013. US Census Bureau (2013)[Zipcode and GPS coordinates]. http://www.census.gov/geo/maps-data/data/gazetteer.html

[2] 2018. NYC OpenData 2018. NYC EMS Incident Dispatch Data. https://data.cityofnewyork.us/Public-Safety/EMS-Incident-Dispatch-Data/76xm-jjuj

[3] Oded Berman, Zvi Drezner, and Dmitry Krass. 2009. Cooperative cover location problems: The planar case. *IIE Transactions* 42, 3 (2009), 232–246. https://doi.org/10.1080/07408170903394355 arXiv:https://doi.org/10.1080/07408170903394355

[4] R. Church and C. Revelle. 1974. The maximal covering location problem. *Papers of the Regional Science Association* 32 (1974), 101–118.

[5] Juan Carlos Dibene, Yazmin Maldonado, Carlos Vera, Mauricio de Oliveira, Leonardo Trujillo, and Oliver Schütze. 2017. Optimizing the location of ambulances in Tijuana, Mexico. *Computers in biology and medicine* 80 (2017), 107–115.

[6] David J. Eaton, Mark S. Daskin, Dennis Simmons, Bill Bulloch, and Glen Jansma. 1985. Determining Emergency Medical Service Vehicle Deployment in Austin, Texas. *Interfaces* 15, 1 (1985), 96âĂŞ108. https://doi.org/10.1287/inte.15.1.96

[7] Michel Gendreau, Gilbert Laporte, and Frédéric Semet. 1997. Solving an ambulance location model by tabu search. *Location science* 5, 2 (1997), 75–88.

[8] Michel Gendreau, Gilbert Laporte, and Frederic Semet. 2001. A dynamic model and parallel tabu search heuristic for real-time ambulance relocation. *Parallel Comput.* 27, 12 (2001), 1641 – 1653. https://doi.org/10.1016/S0167-8191(01)00103-X Applications of parallel computing in transportation.

[9] CM Gibson. 2001. Time Is Myocardium and Time Is Outcomes. *American Heart Association* (2001).

[10] S. L. Hakimi. 1964. Optimum Locations of Switching Centers and the Absolute Centers and Medians of a Graph. *Oper. Res.* 12, 3 (June 1964), 450–459. https://doi.org/10.1287/opre.12.3.450

[11] Averbakh Igor, Berman Oded, Krass Dmitry, Kalcsics Jarg, and Nickel Stefan. 2014. Cooperative covering problems on networks. *Networks* 63, 4 (2014), 334–349. https://doi.org/10.1002/net.21549 arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/net.21549

[12] Gilbert Laporte, François V Louveaux, Frédéric Semet, and Arnaud Thirion. 2009. Application of the double standard model for ambulance location. In *Innovations in distribution logistics*. Springer, 235–249.

[13] Jeffrey L Saver. 2006. Time is brain quantified. *Stroke* 37, 1 (2006), 263–266.

[14] Verena Schmid and Karl F Doerner. 2010. Ambulance location and relocation problems with time-dependent travel times. *European journal of operational research* 207, 3 (2010), 1293–1303.

[15] Masoud Swalehe and Semra Gunay Aktas. 2016. Dynamic Ambulance Deployment to Reduce Ambulance Response Times using Geographic Information Systems: A Case Study of Odunpazari District of Eskisehir Province, Turkey. *Procedia Environmental Sciences* 36 (2016), 199–206.

[16] TC van Barneveld, Sandjai Bhulai, and Robert D van der Mei. 2016. The effect of ambulance relocations on the performance of ambulance service providers. *European Journal of Operational Research* 252, 1 (2016), 257–269.