Finding Optimum Fire Station Locations in Manhattan District

Yaoming Liu[1,a,*,†], Yihao Li[1,b,†]

[1]*Shenzhen College of International Education, Shenzhen, 518043, China*
a. s22185.Liu@stu.scie.com.cn, b. s22427.Li@stu.scie.com.cn
**corresponding author*
[†]*These authors contributed equally to this work and should be considered co-first authors.*

Abstract: Location science has a rich historical background, with the sub-category of locating emergency facilities being especially crucial to ensure the safety of people. With fire emergencies happening frequently in households, identifying the optimal placement of fire stations is a critical decision. Other research papers have used the GIS system or models that are complicated or not as practical. Response time, efficiency, and total cost are the various variables considered and analyzed in the model. Fire stations must be located carefully to ensure they are positioned effectively. Hence, research, data collection, and data analysis are especially important. Data are converted into a weighting system created to compare the importance of different areas to decide whether a fire station should be placed next to them. This project introduces an optimization model for determining the optimal location of fire stations in the Manhattan district using a Python program.

Keywords: Fire Station Locations, Location Science, Python Program.

1. Introduction

The placement of fire stations is a key issue in urban planning and public safety. In densely populated and urbanized areas like New York City, the demand for an efficient and well-distributed fire station network has become more obvious. Especially in Manhattan, it is one of the most densely populated areas in the world. This poses a significant challenge for emergency services. According to the US Census Bureau, by 2020, the population of Manhattan has reached 1.63 million, living in 22.7 square miles. This means that there is an average of 72,000 people per square mile [1]. This density increases the likelihood of fire and the shorter time required for rapid response. At the same time, the complexity and diverse land use in Manhattan, from residential and commercial areas to industrial and park areas, further highlights the need for a strategic approach to relocating fire stations. Because each type of land use has different risk situations and emergency needs. In addition, traffic congestion in the Manhattan area is notorious, which could significantly delay emergency response times. Efficient placement of fire stations can help alleviate delays by ensuring their location will avoid the most severe traffic bottlenecks. According to data from the New York City Fire Department, there were over 1.8 million emergency responses in the city in 2022 [2]. A large part of it occurred in the Manhattan area. The average response time for fires in Manhattan is 4 minutes and 45 seconds [3]. However, in densely populated communities with heavy traffic, response times may be longer.

We believe that this research is necessary. Because this will enhance public safety, faster response times can significantly reduce damage caused by fires and other emergency situations, saving lives and property. Meanwhile, this will make more efficient use of resources. By strategically deploying fire stations, the municipal government can better utilize limited resources such as personnel and land allocation. Ensure that devices are deployed where they are most needed. In addition, this will also improve fairness and ensure that all areas, especially high-risk communities and areas with insufficient services, have enough fire-fighting facilities. It is crucial for public safety. Finally, this study will adapt to urban changes. With the development of cities, the risk landscape will change, and continuous research and optimization ensure that fire services can effectively adapt to these changes.

The main objective of this study is to optimize the location of fire stations to minimize response time and ensure corresponding coverage in different regions.

2. Methods

2.1. Objective function

The primary objective of this research is to optimize the distribution of fire stations in the Manhattan district. By considering the weights of different types of areas. The size of the weight represents the importance of the region, and the larger the weight, the more important it is. Consequently, the areas with higher weights should be closer to fire stations and require less response time when a fire occurs. And how to calculate the weight will be mentioned later. So first of all, we would like to solve the distance problem, trying to find the best way to calculate the distances between fire stations and each demand point. The buildings in Manhattan are mainly distributed in a grid pattern. Due to the grid-like distribution of cities, we do not consider the geometric distance between two points, as there are usually no roads crossing buildings, which means that fire trucks can only travel horizontally and vertically.

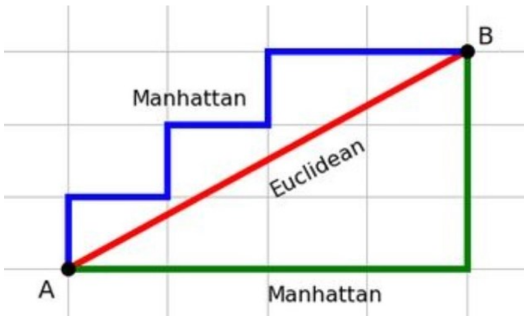

Figure 1: Description of Manhattan distance [4]

Therefore, Manhattan distance is used. As shown in Figure 1, Manhattan distance is the distance between two points in the north or south direction plus the distance in the east or west direction, which can be represented by both the green line and the blue line as well [5]. Once the problem with calculating distance is solved, the trial of finding the optimum location of fire stations is started. Writing down the objective function would help with setting up the foundation of the model.

The p-center model [6] gave us an idea. This model is crucial in ensuring that the placement of fire stations minimizes the maximum distance that any demand point must reach the nearest station. The following is the usage method of the p-center model. We will also explain how we learned from the model and write our objective function.

$$Min \;\; C$$

This equation aims to minimize C, which is the key requirement in the P-center model. C is the maximum distance from any demand point to the nearest fire station, ensuring that the farthest demand point is still within a reasonable distance from the fire station. But in our objective function, we didn't take into account the distance from the demand point to the nearest fire station. We tried to minimize the sum of distances from the fire stations to each demand point.

Assignment Constraints:

$$\sum_{j=1}^{n} Y_{ij} = 1,\ i = 1,2,\dots,m$$

This constraint ensures that each demand point i, i is assigned to exactly one fire station. It ensures that no demand points are missed and every demand point can be serviced.

Facility Location Constraints: The relationship between assignment and facility location is defined

$$Y_{ij} - X_{ij} \le 0,\ i = 1,2,\dots,m; j = 1,2,\dots,n$$

This ensures that demand points can only be allocated to existing fire stations, maintaining logical consistency in allocation. But we didn't consider and learn this part in our function.

Facility Count Constraint: The model also determined the exact number of fire stations

$$\sum_{j=1}^{n} X_{ij} = p$$

This limitation is crucial for planning purposes, ensuring that solutions comply with budget and logistical constraints. We also determined the number of fire stations in advance, but we did not use numerical equations to represent it. It was mentioned in the explanation of the objective function later.

Therefore, the P-Center model provides a rigorous framework for determining the optimal location of fire stations. We have introduced weights to prioritize areas with high demands, ensuring effective and fair allocation of emergency services. Then, we write our Objective function.

$$\min \sum_{i\in I} w_i\left(\left|x_i - x_f\right| + \left|y_i - y_f\right|\right)$$

Objective Function Components

Weight: Wi

The weight represents the importance or demand of each specific point in the area. These weights are assigned based on various factors, such as population density, the frequency of fire incidents, or the presence of critical infrastructure. Higher weights indicate regions that require quicker response times due to higher risks or greater demand for fire services. Therefore, in the objective function, weights are used to prioritize minimizing distances to these high-importance areas, ensuring that the fire station is positioned to serve them effectively.

Demand Points: i

Each I represents a demand point, which represents every single square on the map, such as a commercial area or any other location. The set I includes all these points within the area under consideration. The goal is to ensure that the fire station is optimally placed to minimize the weighted sum of distances to all these points.

Fire Station Coordinates: X_f, Y_f

The coordinates represent the locations of the fire stations that we aim to determine. The objective is to find the optimal coordinates such that the sum of the weighted Manhattan distances to all demand points is minimized.

It seems that our equation is for one fire station, but when we want to set up more fire stations, the coordinates of the fire stations will be enumerated and repeatedly entered into this equation. Find the combination with the smallest total in the end, which is the location where multiple fire stations should be set up. The specific enumeration methods will be mentioned in the next section. By minimizing the objective function, we found the optimal location for the fire station to ensure the fastest response to emergency situations, especially in high-weight areas. Using weights can effectively allocate resources. Prioritize areas with high demand or risk, ensure that fire stations are strategically located, and effectively serve the community. This optimization method supports urban planners in making decisions regarding public safety infrastructure. It ensures that fire stations are placed in the areas that have the greatest impact on community safety. By minimizing the weighted Manhattan distance, the average response time to emergency situations can be shortened, potentially saving lives and property.

2.2. Enumeration and Greedy Algorithm

In order to optimize the positions of multiple fire stations, we initially chose to list all possible combinations to provide a comprehensive and accurate solution for determining the optimal location. This method ensures that all fire station locations minimize the total distance to all demand points. Thus, maximizing the efficiency and effectiveness of the fire stations.

2.2.1. Enumerate Every Single Combination

The first step of this method is to systematically list each possible combination of fire station placement. This involves spreading all possible rows and columns in the predefined grid, which is the Manhattan area in a particular problem. Each unit element in the grid corresponds to a potential location of a fire station. To achieve this goal, the process can be visualized as a nested loop, where one row and one column each time are searched. Here is an example of enumerating the locations of two fire stations.

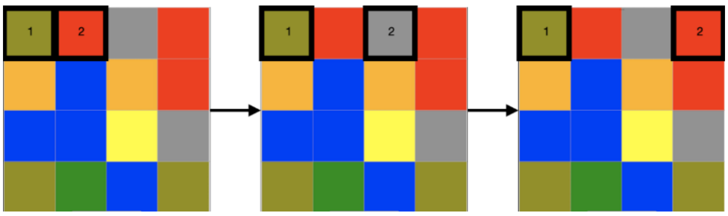

Figure 2: Example of enumeration

As we can see in Figure 2, when we were doing an enumeration of 2 fire stations, we enumerated the combination. The program first fixed one fire station's position and tried all possible positions of another fire station. Then, after finishing the first round (fixed 1 fire station at the first square), the model will jump to the next square (The fixed position will move to the second square and do the same thing as I described in previous words). The different colors represent different types of regions; for example, blue is a residential area. This will be mentioned later in the actual Manhattan district problem. This detailed enumeration is computationally intensive, but it is highly accurate in determining the optimal combination. The advantage of this method is its thoroughness, as it considers all scenarios without approximation. The second step is to introduce the objective function. Substitute the coordinates of each fire station in each combination into the function and find the combination with the smallest total sum. The final step is to determine the locations combination of the fire stations, in order to obtain the minimum sum weighted distance. This combination is considered optimal as it minimizes the average distance that any demand points must travel to reach the fire station while also prioritizing high-demand areas. And I have discovered a small flaw in the objective function here. The 'min' identifier should not be added to the function but rather to the sum. If it is added to the function, it will only be limited to discussing one fire station issue. However, we found that relying on a large number of enumeration methods takes a long time. In the Manhattan area, which is divided into about 300 grids, we calculated that it takes 40 seconds to find the positions of two fire stations. If we add one more fire station, it will be equivalent to multiplying the current required time by 300. Therefore, we introduced the greedy algorithm

2.2.2. Greedy Algorithm

Contrary to detailed enumeration methods, greedy algorithms provide a more efficient and expanded for optimizing fire protection placement[7]. This method builds solutions incrementally, always selecting the next best option based on the current state. Although it may not always guarantee global minimization, greedy algorithms have low computational intensity and are well-suited for large datasets. In simple terms, greedy algorithms work through a series of local optimizations, which means each step will achieve the most direct benefit. This method does not seem like the best option, but when we have enough fire stations. There will be a large number of local optimizations and then combined into a complete distribution of fire stations. The error value will be very small, which can also be proven in our research results about the Manhattan area.

The first step of the greedy algorithm is to determine the optimal initial location of the fire station. It is necessary to first evaluate each potential location and select the location that can best reduce the weighted distance. The calculation of weighted distance still refers to the objective function. This step ensures that the positioning of the initial fire station can maximize its impact on shortening the response time of critical areas. The second step is to place the remaining fire stations in sequence. The program will recalculate the weighted distance of the remaining demand points considering the location of the already placed fire stations and find a solution that can further reduce the weighted distance. The next fire station will be placed in this position. Whenever a new fire station is placed, the weighted distance will be updated. It can be understood that we need to eliminate the weighted distance of the front fire station, preventing the weighted distance from being calculated repeatedly. This iterative placement process will continue until the required number of fire stations are placed. Each step focuses on making the best choice immediately, ensuring that it is the optimal option for the current situation.

2.3. Weight Calculations

It is crucial to calculate the weight of each demand point when determining the optimal location for fire stations. We analyzed population density and the probability of fire occurrence in different regions, and calculated the weights. Firstly, we divided the Manhattan area into many equally sized squares, with actual side lengths of 500m. Therefore, they form a 7-by-43 rectangle representing the entire Manhattan district. Then we use GIS information, shown in Figure 3, to divide the area into five types: residential area, industrial area, commercial area, wood, and landfill. In GIS, different colors are used to indicate different regions. For example, pink represents manufacturing districts, green represents parks and so on.

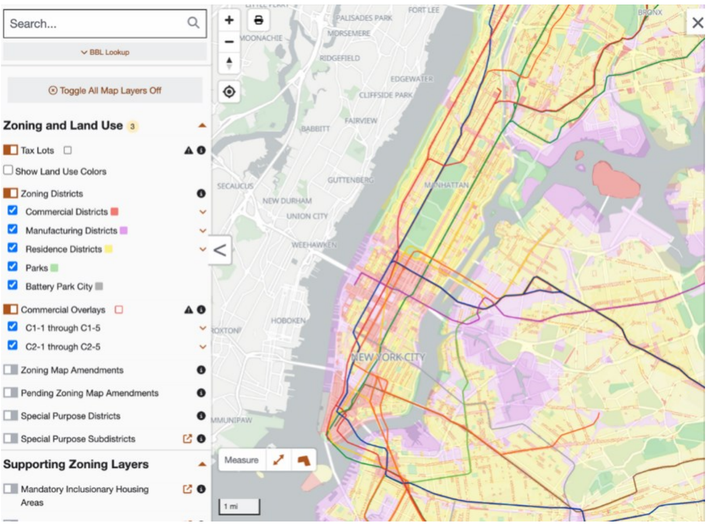

Figure 3: Different types of areas of Manhattan district [8]

Initially, we attempted to use a program to automatically identify the colors in the regions and draw them onto our rectangular grid map. But we soon found that computers had difficulty distinguishing areas where colors overlapped, such as the area below the central park where there were both residential and commercial areas. At the same time, computers also have difficulty distinguishing similar colors, and there are many adjacent light green and bright yellow colors in GIS, which makes our program always recognize errors. Unable to achieve ideal results. Therefore, we switched to manually inputting the colors of each region into the grid map, greatly reducing errors. The result will be shown later, and you can clearly see that the division of different types of regions almost matches the GIS system.

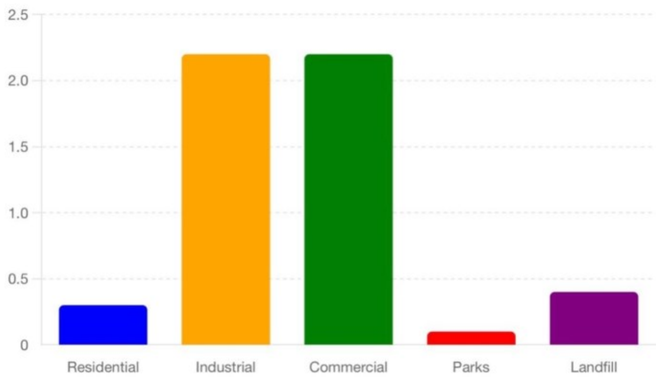

Figure 4: Bar chart of fire emergency probability in different areas

The probability of fire varies in different regions, such as industrial areas, due to the nature of the activities carried out there. The probability of fires occurring in industrial areas is also higher. According to our statistics and review of information, we have found that the probability of a fire in residential areas of Manhattan in 2022 is 0.3% [9], Industrial is 2.2% [10], Commercial is 2.2% [11], Parks or Woods is 0.1% [11], and Landfill is 0.4% [11]. In order to clearly see the difference between the probabilities, Figure 4 is a bar chart that separately records and shows the data of different types of regions, which will also be more convenient for us to mark the weights. The probability of fire directly affects the weight, indicating whether the area is a high-risk area and whether there is a high demand for fire stations.

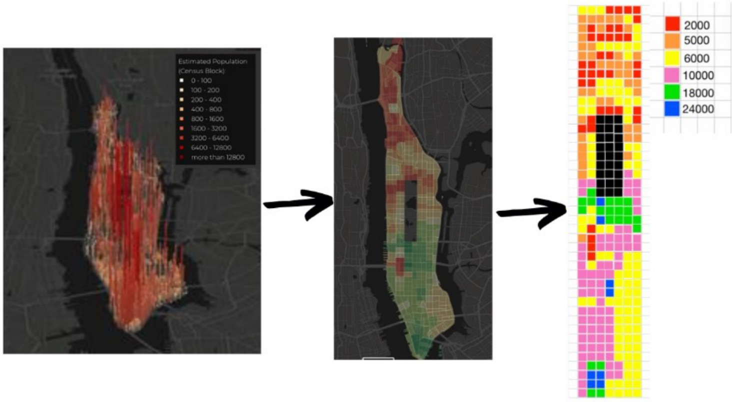

Figure 5: Population distribution of Manhattan district[12] [13]

Secondly, we considered the issue of population distribution and used different color-coded grids to represent the density of different populations. The higher the population density, the more severe the consequences of fires.

At the beginning in Figure 5, we found a 3D image that provided a detailed description of the population distribution, with a vast and accurate amount of information covered in the image. However, this is not conducive to converting the image into a rectangular grid. So, we searched for another plane image. By using the same method, after the program automatically recognized colors, we manually adjusted to reduce errors and finally drew a population distribution map. It is not difficult to find that the population is mainly concentrated in the areas below the central park.

Consequently, we obtained two 7 by 43 grid maps representing different types of regions and population distribution. To calculate the weight of each grid, we introduced a weight calculation formula.

$$Weight = (\Pr o\, bability\ of\ Fire \times 1000) + \left(\frac{Population\ Density}{1000}\right)$$

It is not difficult to find that the value of probability of fire and the numerical value of population data have a significant gap, one in percentage form and the other in large numerical form. So, in order to balance the impact of the two types of data on the weights, we chose to use multiplication by 1000 and division by 1000 processing methods. This processing method not only changes the proportional relationship between the same type of data but also balances the influence of the two types of data. So, each grid has a corresponding weight in this time.

3. Results

The selected fire station locations should satisfy the demand of the Manhattan district. The results, represented in Figure 6, are found according to objectives, which are minimizing response time and total cost. The model considers the weight values of grids when calculating the distance from the fire station to the grid. The multi-objective mathematical model is built to find the best solution and the most appropriate locations for new fire stations.

The result of the mathematical optimization model indicated that the best locations for the fire stations are the areas marked by black squares in the provided grid map. The grid map represents the Manhattan district, with the following colors for different types of areas: residential areas are blue, industrial areas are green, woods are brown, commercial areas are red, and landfills are pink. The optimal locations ensure comprehensive coverage of the Manhattan district, meeting the total demand effectively. Establishing fire stations at these optimal locations will enhance emergency response times and service efficiency across the district.

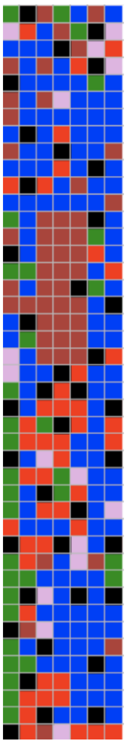

Figure 6: Result of Calculated Optimum Locations of Fire Stations in Manhattan District

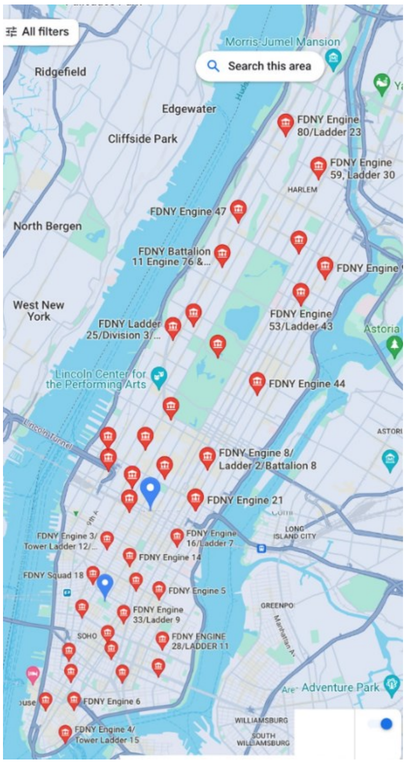

Figure 7: Current Fire station distribution in the Manhattan district [14]

4. Conclusion

Figure 7 is the map of the Manhattan district, with all locations of fire stations marked with red pointers. It is clear that Figure 6 and Figure 7 do not look the same after comparison. To see it more clearly, the map should be split into three parts. Above the park, within/next to the park, and below the park, where the park is the middle area of the map indicated as woods in brown in Figure 6. There are 9 fire stations above the park, 6 fire stations within/next to the park, and 23 fire stations below the park in the calculated results, Whereas there are 5 fire stations, 7 fire stations, and 26 fire stations in the Manhattan district respectively [14].

The differences between our results and the real city could have been caused by several challenges and limitations during the implementation we have faced. First, it was hard to draw the Manhattan map correctly, with the right scale and shape. This made our model less accurate because the map wasn't a perfect representation of the real area. Second, the weights we gave to different factors like population density and land use types (residential, industrial, landfill areas) were not perfectly balanced. This imbalance might have affected our results by making some areas seem more or less important than they actually are or should be in real life. Lastly, the greedy algorithm we used was fast but not the most accurate. It made decisions step-by-step without looking at the whole picture, which sometimes led to less-than-perfect results compared to the more thorough enumeration method. Enumeration considers every possible combination but takes too long to compute, making it impractical for our needs.

Enumeration is a method that evaluates all possible solutions to find the absolute best one. In the context of placing fire stations, enumeration would consider every possible combination of fire station locations to determine which set minimizes response times most effectively. One of the main advantages of enumeration is its accuracy. Enumeration guarantees that the optimal solution is found because it checks every possibility. This comprehensive approach considers the global context, meaning it evaluates how different locations interact with each other to provide the best overall solution. However, enumeration is extremely time-consuming. Because it checks every possible combination, it requires a lot of computational time and power, especially as the number of locations increases. For large-scale problems with many possible locations, enumeration can take an impractically long time, sometimes even years, to produce results. This makes it impractical for real-time decision-making where quick results are essential.

The greedy algorithm, on the other hand, makes a series of local optimizations. It starts by selecting the best option at the current step without considering the entire problem. For placing fire stations, it might start by placing one station in the best possible location and then placing each subsequent station based on the best choice available at that moment. The primary advantage of the greedy algorithm is its speed. The greedy algorithm is much faster because it makes decisions one step at a time and doesn't need to evaluate every possible combination. This quick computation makes it suitable for situations where decisions need to be made quickly, like in emergencies. However, the greedy algorithm focuses on local optimization, meaning it might not find the best overall solution because it only considers immediate benefits rather than the whole picture. By making decisions step-by-step, it can miss out on better solutions that could have been found with a more comprehensive approach, like enumeration.

Despite these challenges, the project shows a useful way to combine GIS with a mathematical model to decide where to place fire stations. The resulting map, though not perfect, gives a good starting point for further improvements. The greedy algorithm, while not as accurate as enumeration, is much faster, making it practical for real-time decisions. This means that in an emergency, city planners can quickly identify potential fire station locations without waiting for a computer to finish long calculations.

For future research, we plan to make the program work better with irregular maps and areas, which will fix the problem of having an imperfectly scaled map. This means that the program will be able to handle more complex shapes and sizes, making it more accurate and useful for different cities and neighborhoods. We also want to consider actual roads and routes, not just the straight-line (Manhattan) distance between points. This will make our model more realistic, as it will reflect the true travel time and routes that fire trucks would use in real emergencies. While enumeration is very accurate, it takes a long time, sometimes even years, when calculating multiple fire stations. So, we use the greedy algorithm, which calculates quickly (within seconds) by testing fire station locations one by one. This balance between speed and accuracy makes the greedy algorithm a valuable tool. It allows us to find good solutions quickly, even if they are not perfect, which is crucial in emergency situations where every second counts.

This decision support system is a prototype for more complex solutions in placing emergency facilities and solving other real-life spatial problems. It shows how technology can help make important decisions that keep people safe. By continuing to improve the model and address its current limitations, we can create even better tools for disaster management and urban planning. This study is a step towards better planning and response strategies, demonstrating the potential of integrating GIS and mathematical models to solve complex problems in urban environments.

Acknowledgement

Yaoming Liu and Yihao Li contributed equally to this work and should be considered co-first authors.

References

[1] Manhattan. (2024, August 22). Wikipedia. https://en.wikipedia.org/wiki/Manhattan

[2] Emergency Response Incidents | NYC Open Data. (2024, August 21). https://data.cityofnewyork.us/Public-Safety/Emergency-Response-Incidents/pasr-j7fb/about_data

[3] New dispatch protocol speeds FDNY response times. (2019, September 18). FireRescue1. https://www.firerescue1.com/fire-department-management/articles/new-dispatch-protocol-speeds-fdny-response-times-KQmbkNX1onSCzAWR/

[4] De Wet, R. (2024, January 18). Manhattan Distance Calculator. Omni Calculator. https://www.omnicalculator.com/math/manhattan-distance

[5] What Is The Manhattan Distance? (n.d.). https://algodaily.com/lessons/what-is-the-manhattan-distance

[6] Tao, L., Cui, Y., Xu, Y., Chen, Z., Guo, H., Huang, B., & Xie, Z. (2022). Location Optimization of Urban Fire Stations Considering the Backup Coverage. International Journal of Environmental Research and Public Health, 20(1), 627. https://doi.org/10.3390/ijerph20010627

[7] Moore, K., Khim, J., & Ross, E. (n.d.). Greedy Algorithms. Brilliant. https://brilliant.org/wiki/greedy-algorithm/

[8] New York City's Zoning & Land Use Map. NYC Planning ZoLa. (n.d.). https://zola.planning.nyc.gov/

[9] Katz, L., & Marr, T. (2022, September 8). America Is Increasingly Building Homes in Disaster-Prone Areas. Redfin Real Estate News. https://www.redfin.com/news/homes-built-disaster-prone-areas/

[10] The Voss Law Firm. (n.d.). Manufacturing and Industrial Fires Responsible for $1 Billion in Property Damage Annually. Fire Statistics for Industrial and Manufacturing Properties. https://www.vosslawfirm.com/blog/fire-statistics-for-industrial-and-manufacturing-properties.cfm

[11] NFPA. (n.d.). Fires by occupancy or Property Type. NFPA report - Fires by occupancy or Property Type. https://www.nfpa.org/education-and-research/research/nfpa-research/fire-statistical-reports/fires-by-occupancy-or-property-type#

[12] Schulz, D., & Schulz, D. (2023b, August 2). See Manhattan's population pulse over 24 hours | 6sqft. 6sqft | NYC Real Estate News and Information. https://www.6sqft.com/see-nycs-population-pulse-over-24-hours/

[13] (N.d.). Manhattan Population Explorer. map. Retrieved from https://manpopex.us/.

[14] (N.d.). Google Maps. map. Retrieved from https://www.google.com/maps/search/fire+station/@40.7505616, -74.0262637, 13z/data=!3m1!4b1!4m8!2m7!3m6!1sfire+station!2sWashington+Square+Park, +Washington+Square, +New+York, +NY+10012, +United+States!3s0x89c25990e23d7e8d:0xfa615edfd1b67e18!4m2!1d-73.997332!2d40.7308838?entry=ttu.