# A Matheuristic Approach for Finding Effective Base Locations and Team Configurations for North West Air Ambulance

Burak Boyacı
Lancaster University, Centre for Transport and Logistics (CENTRAL)
Lancaster, United Kingdom
b.boyaci@lancaster.ac.uk

Muhammad Ali Nayeem*
Lancaster University, Department of Management Science
Lancaster, United Kingdom
m.nayeem@lancaster.ac.uk

Ahmed Kheiri
Lancaster University, Centre for Transport and Logistics (CENTRAL)
Lancaster, United Kingdom
a.kheiri@lancaster.ac.uk

## ABSTRACT

North West Air Ambulance (NWAA) provides helicopter emergency medical service to the Northwest of England. Their three healthcare teams provide their service from two bases with three helicopters. They face some research questions to understand the impact of the air ambulance base locations and the healthcare teams' composition on their services. This paper aims to address those questions by modelling their operations into a location-related decision problem. Then we developed a matheuristic approach to solving the model and generated many realistic instances from historical data to validate our proposed approach's robustness. With the help of our experimental results, we examine the effect of adjusting air ambulance base locations as well as team configurations on the service quality measures to answer the research questions.

## CCS CONCEPTS

• **Computing methodologies → Randomized search**.

## KEYWORDS

Real-world Application, Facility Location Problem, Matheuristics

**ACM Reference Format:**
Burak Boyacı, Muhammad Ali Nayeem, and Ahmed Kheiri. 2021. A Matheuristic Approach for Finding Effective Base Locations and Team Configurations for North West Air Ambulance. In *2021 Genetic and Evolutionary Computation Conference Companion (GECCO '21 Companion), July 10–14, 2021, Lille, France.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3449726.3463151

## 1 INTRODUCTION

The focus of this paper is a real-life research problem originated at and posed by an organisation called North West Air Ambulance (NWAA). Funded by charitable donations, NWAA are dedicated to providing emergency medical services in the Northwest of England in the Greater Manchester, Lancashire, and Cumbria regions. They aim to take advanced healthcare to the spot to improve chances of survival and reduce the risk of long-term injury. The organisation currently consists of six vehicle assets, three air ambulances (i.e., helicopters), and three rapid response vehicles, and can serve more than 2,000 incidents per year. There are three advanced healthcare teams and two base sites in NWAA; two healthcare teams are located near Manchester at Barton and one at Blackpool. The team, based in Barton, is composed of a highly trained doctor and specialist paramedics. The other two teams are based in Blackpool and Barton and are composed of only paramedics. Although the paramedics working for NWAA are more experienced than the usual paramedics, their emergency response capabilities are limited compared to a medical doctor. Each team may be assigned to a single helicopter, so no more than three NWAA assets are in use at any given time instance. A list of the existing helicopters and the associated team is provided below:

- Helicopter H08 is Paramedic crewed based at Blackpool and supported by HX002 Rapid Response Vehicle
- Helicopter H72 is Doctor crewed based at Barton and supported by HX001 Rapid Response Vehicle
- Helicopter H75 is Paramedic crewed based at Barton and supported by HX003 Rapid Response Vehicle

As a charity, NWAA would like to ensure that the donated money is used appropriately. Therefore, NWAA are keen to understand the impact of the air ambulance base locations and the healthcare teams' composition on their services. In this paper, we leverage the well-studied domain of facility location models and aim to address the following research questions posed by NWAA:

(1) How to identify the best candidate base locations and compare their performances with the existing base locations? How does deploying an additional base improve the quality of service?
(2) How to identify the effect of the teams' capabilities on the efficiency of the operations? How does employing more doctors improve the quality of service?

### 1.1 Facility Location Models for Healthcare

Facility location models have become highly relevant in enhancing healthcare services/systems' performance as they help to attain crucial design objectives of minimising operating costs or maximising social benefits. Furthermore, they can be used to examine service usability issues, evaluate the effectiveness of existing location based decisions, and suggest better alternatives to the current solutions [16]. As a result, the diverse application of such

*Muhammad Ali Nayeem is currently a PhD student at Bangladesh University of Engineering and Technology, Bangladesh.


*GECCO '21 Companion, July 10–14, 2021, Lille, France*

ACM ISBN 978-1-4503-8351-6/21/07…$15.00
https://doi.org/10.1145/3449726.3463151

models is prevalent in the literature (e.g., inclusion of explicit geographical conditions in healthcare planning [10], determining locations of ambulances [17], trauma treatment resources [6], organ transplantation facilities [21], placement of blood banks [12], emergency medical service designs [4], preventive healthcare network design [23], diagnostic test laboratories [19], design of hierarchical health service network [22], layout planning of large and complex hospitals [11], optimisation of helicopter rescue operation [13] etc.). Notably, due to computational intractability, previous studies usually tackled real-world healthcare facility problems through heuristic/metaheuristic approaches. For a survey of healthcare facility location problems, readers are referred to [1]. Besides facility location models, there are some studies [7, 9, 14, 15] on the simulation/investigation of different aspects of helicopter based emergency medical services but we find them to be out of the scope of this paper.

### 1.2 Our Contributions

This study makes the following key contributions. Firstly, we model the helicopter operations of NWAA as a capacitated facility location-allocation variant (Section 2). Our model considers not only base locations but also the number and the compositions of the teams at each base. Then, we construct realistic problem instances that mimic the NWAA activity pattern/distribution by leveraging three historical databases (Section 3). Furthermore, to effectively solve such practical instances, we develop a matheuristic approach by embedding an exact method inside a hill-climbing structure (Section 3) and examine its robustness (Section 4). It enables us to examine the model outputs while varying the input parameters in accordance with the research questions (Section 4) and offer our findings/recommendations (Section 5). Thus this paper combines different aspects of the well-studied facility location models to perform a cost-effectiveness analysis of the model that suits the NWAA operations that is not directly comparable to previous studies on healthcare facility problems.

## 2 PROBLEM DESCRIPTION AND MATHEMATICAL FORMULATION

From historical records, we find that NWAA rarely rejects a call because of serving another operation. Moreover, the helicopters used by NWAA can reach speeds of 150mph, allowing them to reach any incident from a base located in the region's central part within 30 minutes. Because of these reasons, we aim to find base locations that will minimise the average flying distance without considering the specific time (or time window) to serve each incident. We can formulate this problem as a variant of the capacitated facility location-allocation problem. To be more exact, it is a special case of the Capacitated Multi-facility Weber Problem (CMWP). In CMWPs, the aim is to locate a predefined number of facilities in an Euclidean space that will satisfy customers' demand while minimising the total transportation cost. In our problem, we consider each incident (i.e., customer) has unit demand, and the capacity of the bases (i.e., facilities) depends on the number of teams located at each base. Besides, since some incidents require advanced skills, a subset of incidents can only be served by a team with a doctor. As a result, we assume there are two capacities for each base: the number of incidents that require a doctor and the total number of incidents. We also assume each base can serve up to certain percentage of the total incidents. This makes the problem slightly different than the classical CMWPs since the total supply exceeds the total demand. For the given indices, parameters, and variables below, we develop the following mathematical model.

### Sets and Indices

| | |
|---|---|
| $i \in I$ | bases |
| $j \in J$ | incidents |
| $J' \subset J$ | incidents requiring teams with a doctor |

### Parameters

| | |
|---|---|
| $\mathbf{a}_j = (a_{j1}, a_{j2})$ | coordinates of incident $j$ |
| $s_i/s_i'$ | total / {doctor requiring} number of incidents that can be served by base $i$ |
| $d(i, j)$ | distance between base $i$ and incident $j$ |

### Variables

| | |
|---|---|
| $\mathbf{x}_i = (x_{i1}, x_{i2})$ | coordinates of base $i$ |
| $w_{ij}$ | 1 if incident $j$ is assigned to base $i$; 0 otherwise |

### Mathematical Modelling

$$\min \sum_{i \in I} \sum_{j \in J} w_{ij} d\left(\mathbf{x}_i, \mathbf{a}_j\right) \tag{1}$$

$$\text{s.t.} \sum_{i \in I} w_{ij} = 1 \qquad \forall j \in J \tag{2}$$

$$\sum_{j \in J} w_{ij} \le s_i \qquad \forall i \in I \tag{3}$$

$$\sum_{j \in J'} w_{ij} \le s_i' \qquad \forall i \in I \tag{4}$$

$$w_{ij} \in \{0, 1\} \qquad \forall i \in I; j \in J \tag{5}$$

Here, objective function in equation 1 minimises the total distance between the incidents and the bases that serve the incidents. Constraint set 2 force that every incident is served by one and only one base. Constraint sets 3 and 4 ensure that the number of all incidents and incidents requiring a doctor served by each base do not exceed the capacities of the bases, respectively. These capacity constraints ensure all teams' participation, which eventually helps to reserve the doctors for severe cases. Constraint set 5 define the domain of variables $w_{ij}$. It should be noted that, although the decision variables used in this model are the base locations and allocations of each incident to a particular base, in the context of this paper, we are only concerned with the base locations.

In this problem, for $d\left(\mathbf{x}_i, \mathbf{a}_j\right)$, we use an Euclidean distance function since the helicopters can fly directly to the incident locations. However, when the distance function is Euclidean, the objective function of this model is neither concave nor convex [8]. Although a few methods solved this problem exactly (see [3, 18]) or approximately (see [2, 5]), the size of the problems that the algorithms proposed in these papers can solve are quite limited compared to our instances. To have not necessarily optimal but good solutions, we look for heuristic approaches and decide to solve the location and allocation problems iteratively similar to the proposed approach of Cooper [7]. In this approach, we start from a set of random base locations. When we fix the location variables, $x_i$, the problem becomes a variant of the network flow problem with unit flows. We solve this problem exactly with a mathematical solver. The solution to this problem gives us the allocations of the incidents to each base. Fixing the allocation variables, $w_{ij}$, decomposes the problem to (single facility) Weber Problems for each base. We utilise Weiszfeld's procedure [20] to find locations minimising the total distance between each base and incidents assigned to them. We repeat this algorithm multiple times, starting from random initial locations to avoid getting caught up in locally optimal solutions. A more detailed description of this algorithm and its pseudo-code is given in Section 3.

## 3 METHODOLOGY

In this section, we develop an efficient approach for solving the problem formulated in the previous section by combining heuristic and exact techniques. In addition, we discuss our steps involved in the generation of real-world instances from historical databases.

### 3.1 Matheuristic Approach

It is not possible to obtain optimal solutions using an exact method for the formulation presented in Section 2 on realistic instances within a reasonable time. So we devise an efficient matheuristic approach (Algorithm 1 and Figure 1) by adapting Cooper's iterative location-allocation algorithm [7, 8] that acts like a hill-climber from a high-level view. It starts with random base locations and the final solution depends on the initial locations. Before entering the main loop, we ensure the feasibility of the initial solution (line 5) by checking whether each base has been assigned to serve at least one incident from the current allocation. The algorithm alternates between location and allocation steps until a local optimal solution is found. In the location step (line 13), we update the location of each candidate base by the geometric median of the locations of incidents currently allocated to that base (line 3 of Algorithm 2) following Weiszfeld's iterative procedure [20]. However, in the allocation step (line 14), unlike the Cooper's algorithm which uses only one type of demand, we consider incidents that require an intervention of a team with a doctor and all incidents separately. We solve this problem by using a mathematical model given in Section 2 by fixing variables $x_i$ to current candidate base locations. Thus we select the optimal allocation within a sub-portion of the search space to maintain the time-efficiency. The robustness of this method in regards to withholding variation in the problem instances is examined in Section 4.3.1.

**Algorithm 1** Matheuristic Approach

**Input:** *maxIter*: the maximum number of iteration; *nBases*: the number of bases; *teamConfig*: the team configuration; *probInst*: the instance to be solved

1: *model* ← create a model of Section 2 with *nBases*, *teamConfig* and *probInst*
2: **repeat**
3: *curLoc* ← randomly initialise *nBases* locations
4: *curAlloc* ← fix the location related decision variables **x** to *curLoc* and then solve *model* by an exact method to get the allocation
5: **until** *curLoc* and *curAlloc* constitute a feasible solution
6: *bestLoc* ← *curLoc* /*stores the best locations and returned as output*/
7: *bestAlloc* ← *curAlloc* /*stores the best allocation and returned as output*/
8: *bestScore* ← evaluate the objective function (eqn. 1) using *bestLoc* and *bestAlloc* /*stores the best score and returned as output*/
9: *curIter* ← 0 /*iteration counter*/
10: **repeat**
11: *curIter* ← *curIter* + 1
12: *prevAlloc* ← *curAlloc*
13: *curLoc* ← Update-Base-Locations(*curLoc*, *curAlloc*) /*the location step (Algorithm 2)*/
14: *curAlloc* ← fix the location related decision variables **x** to *curLoc* and then solve *model* by an exact method to get the allocation /*the allocation step*/
15: *curScore* ← evaluate the objective function using *curLoc* and *curAlloc*
16: **if** *curScore* is better than *bestScore* **then**
17: *bestScore* ← *curScore*
18: *bestLoc* ← *curLoc*
19: *bestAlloc* ← *curAlloc*
20: **end if**
21: **until** *curIter* > *maxIter* or *curAlloc* has not been changed w.r.t. *prevAlloc*
22: **return** *bestLoc*, *bestAlloc*, *bestScore*

**Algorithm 2** Update-Base-Locations

**Input:** *curLoc*: candidate base locations; *curAlloc*: incidents currently allocated to each base in *curLoc*

1: **for** each base *b* **do**
2: *P* ← a set of coordinates of incidents from *curAlloc* currently allocated to *b*
3: Update the location of *b* in *curLoc* to be the geometric median of all coordinates in *P*
4: **end for**
5: **return** *curLoc*

### 3.2 Construction of Datasets

Three historical databases as reported in Table 1 are utilised in this paper. We use HEMS database to understand the nature of incidents followed by NWAA and sample similar incidents from 999 and TARN databases to prepare two datasets. We aim to obtain robust solutions that can withstand variation in data, variables or assumptions. To accomplish this, under each dataset, we construct 100 different problem instances of 4,000-5,000 incidents chosen randomly following the same sampling criterion. Thus we get 100 different solutions for each dataset which enables us to examine the robustness of our methodology. Below we discuss our approach to generate the two datasets.

### Dataset1

It includes daytime (from 8:00 AM to 20:00 PM) incidents from 999 database that are not necessarily served by NWAA but can meet NWAA dispatch criteria. It follows the similar distribution

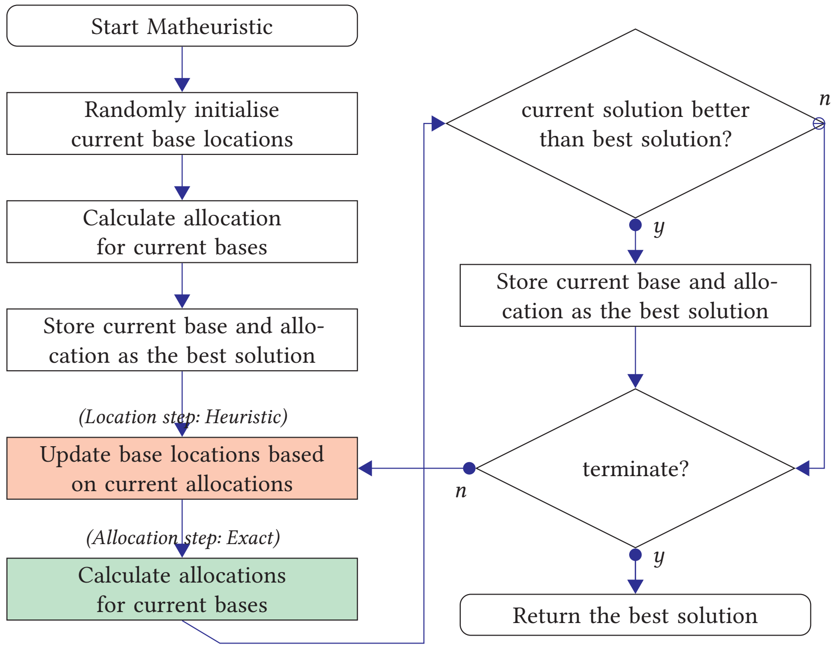


**Figure 1: A high-level flowchart of our matheuristic approach.**

**Table 1: Databases used for construction of problem instances**

| Name | #Entries | Content info. |
|---|---|---|
| HEMS | 4,126 | Calls the NWAA attended / set off to |
| 999 | 1,528,279 | 999 calls from the north west UK |
| TARN | 15,051 | Data on moderately and severely injured patients |

of the HEMS database in terms of the call category, AMPDS code reporting, and healthcare teams (Figure 2). We use the roulette-wheel random selection mechanism to select call category, AMPDS code, and healthcare team for an incident following this distribution. Moreover, to follow the HEMS incidents' spatial distribution (Figure 3) on the map, we split the map area under consideration into several square regions of similar sizes and ensure the inclusion of 999 incidents from each region, maintaining a similar percentage with HEMS. Each square region's size is a design parameter that is tuned to determine the best possible outcome. Following the location distribution will result in a dataset that reflects the general characteristics of the locations served by NWAA, avoiding the clustering of incidents on highly populated cities, such as Manchester and Liverpool.

## Dataset2

Here the goal is to construct a dataset that includes the TARN missions NWAA could have attended. This dataset includes daytime incidents and follows a similar distribution of the HEMS database in only healthcare teams as the TARN database does not have call category and AMPDS code fields. Also, the locations are processed in a similar way as in Dataset1. Moreover, we only consider the most

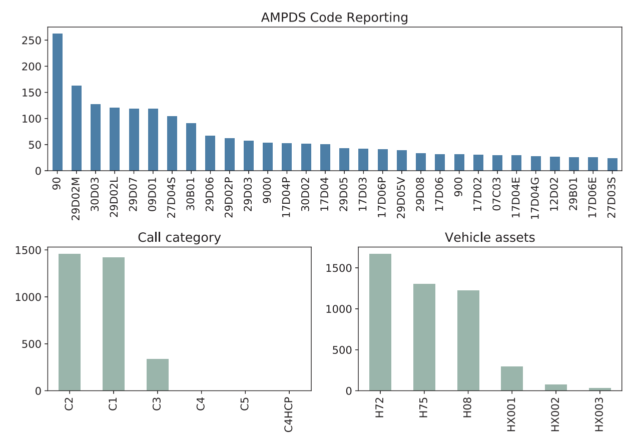


**Figure 2: Distribution of top 30 AMPDS code reporting, call category and vehicle assets from HEMS database. The vehicle assets H72 and HX001 represent healthcare team with a doctor**

severe incidents as NWAA considers severity as the number one priority. This can be identified by looking at the patients with an ISS 'Injury Severity Score' score of 15 and above, which constitutes a major trauma.

# 4 EMPIRICAL STUDY

In this section, we discuss different aspects of our experimentation with our matheuristic approach conducted on two datasets, each

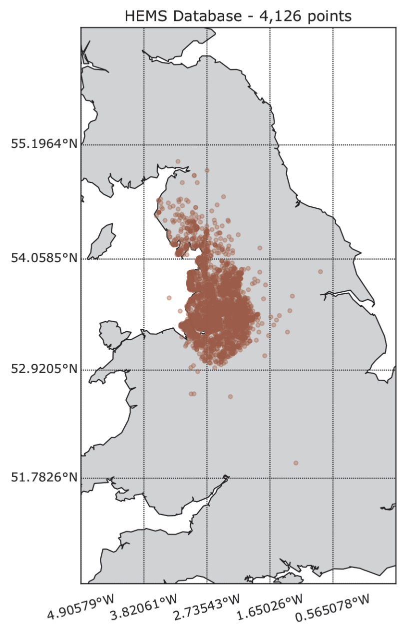


**Figure 3: Locations of incidents from HEMS database**

having 100 problem instances. We implemented our methods using C# and used IBM ILOG CPLEX 12.10 as the mathematical solver.

## 4.1 Experimental Settings

Our analysis considers base locations and the number and configurations of the teams at each base. We assumed that there are three teams available, and we distribute them to 1, 2, or 3 different airbases. Also, we consider different team configurations varying the number of doctors in each team. Thus, we ended up with seven distinct team configurations provided in Table 2. In these team configurations, each base may contain, a number of teams each with a doctor (**D**, **DD**, **DDD**, number of **D** shows the number of teams), a team with only paramedics (**P**) and two teams one with a doctor and another without a doctor (**DP**). Our model aims to minimise total flight distance while serving all incidents of the problem instance.

When we assign incidents to teams, we apply the constraint that if an incident requires a doctor, it can only be served by a team with a doctor; otherwise, the incident can be served by a team with or without a doctor. Besides, we assume every team can serve a fixed percentage (we use 40% in this study) of the incidents at most (constraint 3 in Section 2). This constraint ensures the participation of each team which helps to reserve the doctor team for severe cases. The percentage of the incidents requiring a doctor may be more than the number of teams with a doctor in some instances. To have feasible solutions, in those instances, if required, we relax the number of incidents that can be served by a team with a doctor but force them also to serve only incidents that require a doctor. And for the same reason, we have not created any configuration with only paramedic teams.

For each problem instance, we run our matheuristic method 30 times starting from a different initial solution and take the base locations that achieves the best objective (out of 30) as the solution. Thus for each (dataset, team configuration) combination, we get 100 solutions (i.e., one for each instance). To verify the robustness of our approach, we also evaluate the objective values of a particular solution for the other 99 problem instances which were not used to obtain it. We discuss this in more details in Section 4.3. Table 3 summarises important parameters of our approach.

## 4.2 Performance Measure

In our analysis, we use the total flight distance (nautical miles or NM) and total flight time (hours or h) as the performance metrics. While calculating the flight time, we assume that the helicopter fly with a speed of 134.691 knots (i.e., 155mph). Also, we consider a fixed overhead for every incident spent for take-offs and landings. We calculate total flight distance and total flight time on each problem instance for NWAA's existing base locations and team configuration (i.e., two teams at Barton and one in Blackpool, 1 doctor in Barton only) use these values as a basis for comparing the relative performance of each scenario listed in Table 2.

## 4.3 Experimental Results

Here we present the results of our matheuristic approach across 100 problem instances of the two datasets for different team configurations listed in Table 2 after an examination of its robustness. However, we defer the discussion of our findings/recommendation derived from these results to the next section (i.e., Section 5).

*4.3.1 Validation of Robustness.* To validate our matheuristic approach's robustness, we would like to examine the difference in its performance while varying incidents of the problem instance. To accomplish this goal, randomly split our 100 instances in each dataset into the train (80%) and test (20%) sets similar to a machine learning approach. We obtain 80 sets of base locations by running our method on each train instance and summarise them into one set of base locations by taking the geometric median of the 80 candidates for each base. We refer to this single solution as the *trained model*. We perform these steps for every team configuration. Then we summarise the objective values of the trained model separately for 80 train instances and 20 test instances using one boxplot for each team configuration in Figure 4. To make the comparison more comfortable, we normalise all the objective values for a particular instance obtained across different team configurations using the minimum and maximum. We find each box to be very narrow, implying that the trained model's performances across different instances are quite similar. More importantly, we see that each boxplot of the train results is almost identical to its counterpart in the test results. The same holds for the relative position of each pair of boxplots. In subsequent sections, we see that all the base locations generated by our method across different instances are positioned in close proximity on the map. All these observations indicate our approach's robustness in terms of its flexibility to changes in the incidents.

*4.3.2 Results for Dataset1.* Dataset1 represents daytime 999 calls that follow the distribution of call category, AMPDS code reporting, healthcare teams, and location from the HEMS database. Figure 5

**Table 2: Seven team configurations used in this analysis**

| configuration | base #1 | base #2 | base #3 |
|---|---|---|---|
| **DP\|P** | 2 teams one with a doctor | 1 team | |
| **D\|P\|P** | 1 team with a doctor | 1 team | 1 team |
| **DP\|D** | 2 teams one with a doctor | 1 team with a doctor | |
| **D\|D\|P** | 1 team with a doctor | 1 team with a doctor | 1 team |
| **DDD** | 3 teams each with a doctor | | |
| **DD\|D** | 2 teams each with a doctor | 1 team with a doctor | |
| **D\|D\|D** | 1 team with a doctor | 1 team with a doctor | 1 team with a doctor |

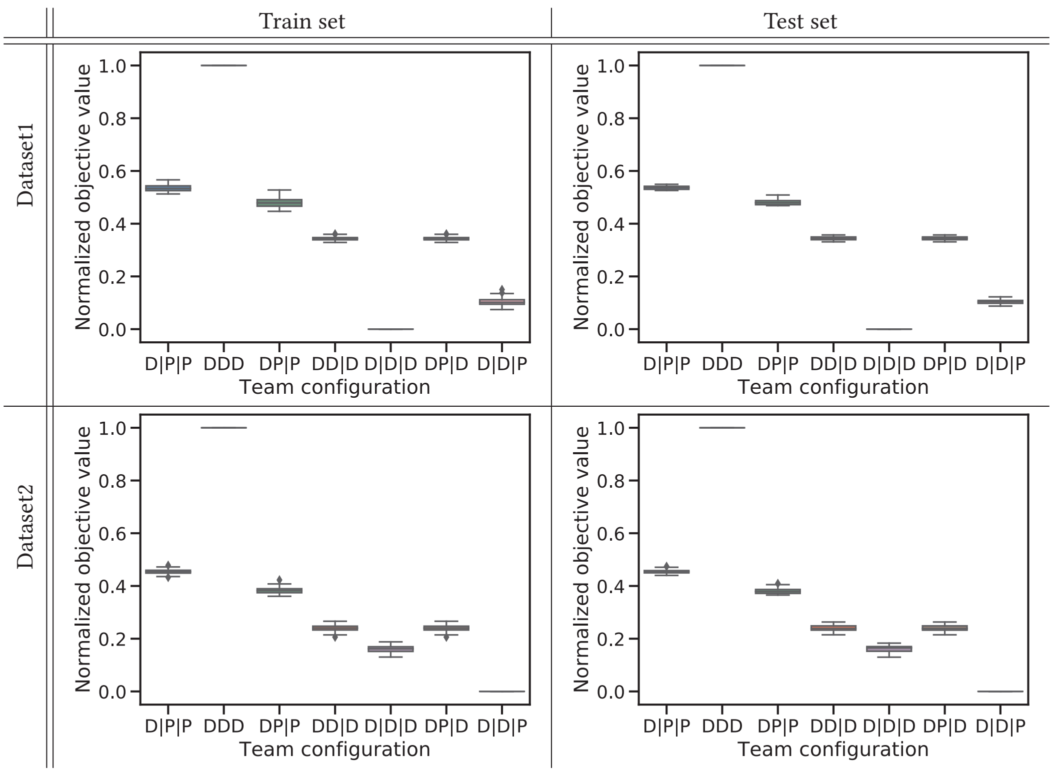


**Figure 4: Performance of the trained model on randomly chosen 80 train and 20 test instances**

**Table 3: Important system parameters**

| Parameter | Value |
|---|---|
| Airbases | 1-3 (see Table 2) |
| Team configurations | 7 (see Table 2) |
| Repeated runs | 30 |
| Max. iterations ($maxIter$) | 100 |
| Max. incidents served by any team ($s_i/s_i'$) | 40% |

summarises our obtained results for Dataset1. In parts (a)-(g), we visually locate all 100 sets of base locations generated by our method on the map for each team configuration. Moreover, part (h) shows the total flight distances and flight time. To calculate the flight distance for a particular team configuration from 100 different solutions, we generate a 100X100 matrix where the cell $(i, j)$ shows the total distance incurred by the solution generated using $i^{th}$ instance but evaluated on the $j^{th}$ instance. Then we transform this matrix into a vector of dimension 100 by averaging all values in the same column. Finally, we average 100 values in the vector to calculate the total flight distance (reported in part (h) of Figure 5) that gives us a robust estimation of the base locations for a particular team configuration considering all instances. We see that the different locations (i.e., one for each instance) identified for a team usually located to each other's proximity implying our approach's robustness. The high population density in Manchester and Liverpool dominates base locations identified in the south, but the base in Liverpool is located to the east of its centre to cover a larger demand area. The

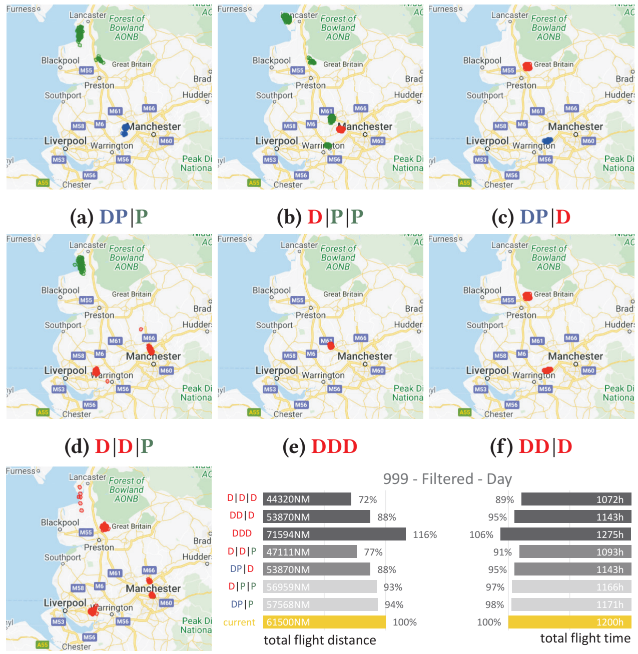


(a) DP|P (b) D|P|P (c) DP|D (d) D|D|P (e) DDD (f) DD|D (g) D|D|D (h) performance measures compared to current configuration

**Figure 5: Visual depiction of all optimised base locations for different team configurations (parts (a)-(g)) obtained on the 100 instances of Dataset1 and their performance comparison (part (h))**

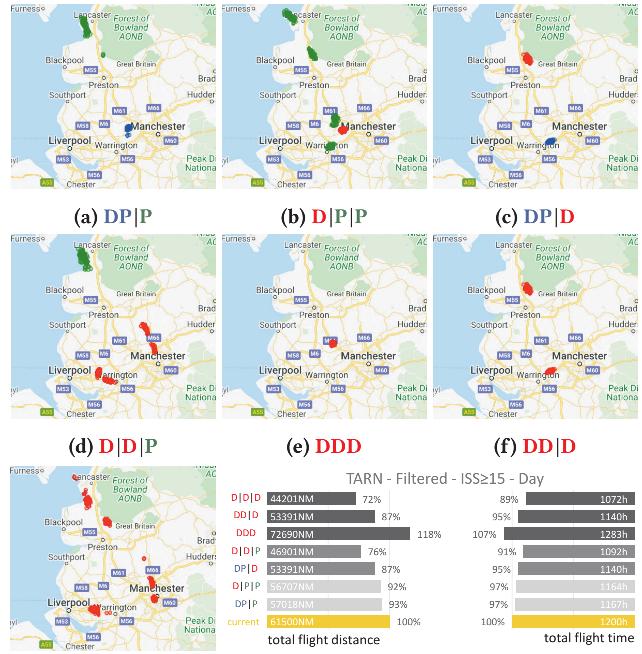


(a) DP|P (b) D|P|P (c) DP|D (d) D|D|P (e) DDD (f) DD|D (g) D|D|D (h) performance measures compared to current configuration

**Figure 6: Visual depiction of all optimised base locations for different team configurations (parts (a)-(g)) obtained on the 100 instances of Dataset2 and their performance comparison (part (h))**

base locations advised in the north are closer to Lancaster, even located in the west of the city.

*4.3.3 Results for Dataset2.* We now present our results for Dataset2 as shown in Figure 6. In this dataset, we include the daytime requests from the TARN database with a minimum ISS score of 15 and follow the distributions of healthcare teams' involvement and incidents' locations from the HEMS database. We do not describe these results in detail as they are very similar to that of Dataset1.

## 5 DISCUSSION

In response to the research questions posed by NWAA, we now discuss our findings and recommendations derived from the experimental results as follows.

- We see that by reconfiguring the location of their north base, NWAA could decrease the total flight distance by 6-7% (around 4000NM/year) and the total flight time by 2-3% (around 30h/year) if they continue to serve as they were in the period we considered.
- The improvement of having a third base when there is only one team with a doctor is tiny (i.e., around 1%) for the total flight distance (which is around 615NM/year) and quite negligible for the total flight time. On the other hand, having a team with another doctor when there are two bases shows huge improvement. In both datasets, having two teams with doctors instead of one when there are two bases decreases the total flight distance by 6% (around 3700NM/year) and the total flight time by 2.5% (around 30h/year) when it is compared to the locations optimised by our approach. Upgrading the current base locations (specifically the north base) and converting one of the paramedics only teams to a team with a doctor could decrease the total flight distance by 12-13% (around 7700NM/year) and the flight time by 5% (around 60h/year).
- Having three bases is beneficial only if there are at least two teams with a doctor. When all the teams have doctors, the third base decreases the total flight distance by 15-16% (around 9500NM/year) and the total flight time by 6% (72h/year). The effect is still substantial but lower when there are two teams with a doctor. If there are only two bases and two teams with a doctor, having a third base decreases the total flight distance by 11% (6700NM/year) and the total flight time by 4% (48h/year).

## 6 CONCLUSION AND FUTURE WORK

In this paper, we addressed some research questions of NWAA regarding the base location and team configuration of their air ambulance operation to maximise their impact. We modelled the problem as a special case of the Capacitated Multi-facility Weber Problem (CMWP) and developed a robust matheuristic approach along with a large number of realistic problem instances. We developed an efficient matheuristic method by combining heuristic and exact techniques in a way that can generate high-quality solutions in a reasonable amount of time. The close proximity of the obtained base locations on the map across 100 different instances indicates our approach's robustness. Finally, upon careful examination of the obtained outputs for different scenarios, we provided our observations/recommendations.

It is important to note that, in our analysis, we use the past incident data. If there is a significant change that affects the incident distribution or count in the region (e.g., Covid-19 Pandemic), these effects could be different. In our analysis, we tried our best and use what is available to us. Having said that, if the demand changes (more likely during a pandemic) but the distribution does not change (less likely during a pandemic), the locations we have identified will still improve the total flight distance and time but less or more depending on the change in demand levels. However, from another angle, after some real-life data are available for a new scenario, our data processing pipeline could be useful in such cases through which we can produce realistic instances to run our method thereon for more realistic output.

We have not analysed how much NWAA could gain/lose if they use specific locations for a given team composition since the number of options could be infinitely many. However, given extra information about the candidate scenarios for base locations and team compositions, we can conduct these additional analyses for NWAA. Currently, we are devising a hyper-heuristic based framework to improve the solution generation process further. We are also working to optimise the operations of rapid response vehicles of NWAA, which is a more complicated task than what we addressed here.

## ACKNOWLEDGMENTS

Authors are grateful to North West Air Ambulance (NWAA) for providing valuable data. Authors would like to thank David Briggs from NWAA for providing valuable feedback. Muhammad Ali Nayeem is supported by the internship scholarship awarded by SPECIES and the ICT Doctoral Fellowship administered by ICT Division, Government of People's Republic of Bangladesh; and received important directions from Dr. M. Sohel Rahman, Bangladesh University of Engineering and Technology.

## REFERENCES


[1] Amir Ahmadi-Javid, Pardis Seyedi, and Siddhartha S Syam. 2017. A survey of healthcare facility location. *Computers & Operations Research* 79 (2017), 223–263.

[2] M Hakan Akyüz, Temel Öncan, and İ Kuban Altınel. 2012. Efficient approximate solution methods for the multi-commodity capacitated multi-facility Weber problem. *Computers & operations research* 39, 2 (2012), 225–237.

[3] M Hakan Akyüz, Temel Öncan, and İ Kuban Altınel. 2019. Branch and bound algorithms for solving the multi-commodity capacitated multi-facility Weber problem. *Annals of Operations Research* 279, 1-2 (2019), 1–42.

[4] Patrizia Beraldi and Maria Elena Bruni. 2009. A probabilistic model applied to emergency service vehicle location. *European Journal of Operational Research* 196, 1 (2009), 323–331.

[5] Burak Boyaci, İ Kuban Altinel, and Necat Aras. 2013. Approximate solution methods for the capacitated multi-facility Weber problem. *IIE Transactions* 45, 1 (2013), 97–120.

[6] Charles C Branas, Ellen J MacKenzie, and Charles S ReVelle. 2000. A trauma resource allocation model for ambulances and hospitals. *Health Services Research* 35, 2 (2000), 489.

[7] Leon Cooper. 1964. Heuristic methods for location-allocation problems. *SIAM review* 6, 1 (1964), 37–53.

[8] Leon Cooper. 1972. The Transportation-Location Problem. *Operations Research* 20, 1 (1972), 94–108. https://doi.org/10.1287/opre.20.1.94

[9] Leigh Curtis, Mark Salmon, and Richard M Lyon. 2017. The impact of helicopter emergency medical service night operations in south East England. *Air medical journal* 36, 6 (2017), 307–310.

[10] Paul Robert Harper, AK Shahani, JE Gallagher, and C Bowie. 2005. Planning health services with explicit geographical considerations: a stochastic location–allocation approach. *Omega* 33, 2 (2005), 141–152.

[11] Stefan Helber, Daniel Böhme, Farid Oucherif, Svenja Lagershausen, and Steffen Kasper. 2016. A hierarchical facility layout planning approach for large and complex hospitals. *Flexible services and manufacturing journal* 28, 1-2 (2016), 5–29.

[12] Armin Jabbarzadeh, Behnam Fahimnia, and Stefan Seuring. 2014. Dynamic supply chain network design for the supply of blood in disasters: A robust model with real world application. *Transportation Research Part E: Logistics and Transportation Review* 70 (2014), 225–244.

[13] Mumtaz Karatas, Nasuh Razi, and Murat M Gunal. 2017. An ILP and simulation model to optimize search and rescue helicopter operations. *Journal of the Operational Research Society* 68, 11 (2017), 1335–1351.

[14] Richard M Lyon, Joe Vernon, Magnus Nelson, Neal Durge, Malcolm Tunnicliff, Leigh Curtis, and Malcolm Q Russell. 2015. The need for a UK helicopter emergency medical service by night: a prospective, simulation study. *Air medical journal* 34, 4 (2015), 195–198.

[15] Lucy Morgan and Roger Brooks. 2020. Strategic Decision Making for the North West Air Ambulance Charity Using Discrete Event Simulation. In *SW20-Operational Research Society-10th Workshop*.

[16] Shams-ur Rahman and David K Smith. 2000. Use of location-allocation models in health service development planning in developing nations.

[17] 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.

[18] Hanif D Sherali, Intesar Al-Loughani, and Shivaram Subramanian. 2002. Global optimization procedures for the capacitated Euclidean and lp distance multifacility location-allocation problems. *Operations Research* 50, 3 (2002), 433–448.

[19] Honora Smith, Daniel Cakebread, Maria Battarra, Ben Shelbourne, Naseem Cassim, and Lindi Coetzee. 2017. Location of a hierarchy of HIV/AIDS test laboratories in an inbound hub network: case study in South Africa. *Journal of the Operational Research Society* 68, 9 (2017), 1068–1081.

[20] Endre Weiszfeld. 1937. Sur le point pour lequel la somme des distances de n points donnés est minimum. *Tohoku Mathematical Journal, First Series* 43 (1937), 355–386.

[21] Behzad Zahiri, Reza Tavakkoli-Moghaddam, Mehrdad Mohammadi, and Payman Jula. 2014. Multi-objective design of an organ transplant network under uncertainty. *Transportation Research Part E: Logistics and Transportation Review* 72 (2014), 101–124.

[22] Naeme Zarrinpoor, Mohammad Saber Fallahnezhad, and Mir Saman Pishvaee. 2018. The design of a reliable and robust hierarchical health service network using an accelerated Benders decomposition algorithm. *European Journal of Operational Research* 265, 3 (2018), 1013–1032.

[23] Yue Zhang, Oded Berman, and Vedat Verter. 2009. Incorporating congestion in preventive healthcare facility network design. *European Journal of Operational Research* 198, 3 (2009), 922–935.