



# GIMA

Geographical Information Management and Applications

## Location allocation of mobile hospitals in a Complex Humanitarian Emergency

*Final report*



Alistair Steward  
[a.p.steward@students.uu.nl](mailto:a.p.steward@students.uu.nl)  
Supervisor: Ellen-Wien Augustijn  
Responsible professor: Raul Zurita-Milla  
February 2020

## Abstract

Complex humanitarian emergencies (CHEs) resulting from conflict are catastrophic for public health, with dramatic increases in traumatic injuries and primary level healthcare needs compounded by a breakdown in existing health infrastructure. Mobile hospitals can support overburdened or damaged health facilities by providing facilities equivalent to first, second and in some instances third level hospitals, but limited work has been carried out on where best to locate them within a dynamic emergency environment. This research attempted to fill the gap by constructing a spatial decision support system (SDSS) for decision makers who are not knowledgeable in geoinformation processes to quickly generate suitable locations for mobile hospitals. Mosul, Iraq, a region that experienced a CHE in 2016 and 2017, was used as a case study to develop the application.

Through a literature review of mobile hospital characteristics, various location criteria were identified. The most important of these were the safety of the hospital patients and staff, the amount of potential patients nearby, and resilient site accessibility because road networks were likely to be unstable during an emergency. These were taken as assessment criteria for the SDSS, which opted for a fast, lightweight genetic algorithm over traditional multi-criteria evaluation processes. The SDSS was broken into three key phases. Analysis of 20 locations supplied by the WHO informed the generation and filtering of a candidate location grid for the SDSS to evaluate. Users were then able to drop candidates in unsafe restricted areas, which varied temporally, before they were fed into a Non-dominated Sorting Genetic Algorithm II for the final phase. This used the mechanics of natural selection to return pairs of candidates with the highest cumulative population number and centrality scores.

To determine the extent that recommended locations changed over time, results were collected by running the SDSS 300 times for four different timesteps and a benchmarking exercise that did not include restricted areas. Analysis of these results demonstrated that solutions varied substantially depending on the timestep, reinforcing the point that locating mobile hospitals is a challenging task for decision makers, and that an SDSS would be a useful asset. A comparison was also carried out with models developed using the same case study, exposing strengths and weaknesses of the SDSS. It was significantly faster than the more conventional GIS-focused model of Amer, Augustijn, van den Bosch and Da Silva Mano (2017) but lacked sophistication when it came to determining service areas and integration with the healthcare network.

Conclusions from the research corroborated previous work in favour of using genetic algorithms for location allocation. Although the final SDSS had several limitations, these are well understood and it has a lot of potential as a foundation for future work. Recommendations include carrying out a formal usability assessment and integration of a mechanism to better cater for location of multiple mobile hospitals simultaneously.

# Contents

<b>1. Introduction</b>	<b>3</b>
1.1 Research context	3
1.2 Research objectives and sub-questions	5
1.3 Scope	6
1.4 Limitations at the outset	7
1.5 Research framework	7
1.6 Structure of this document	9
<b>2. Literature review</b>	<b>10</b>
2.1 Location criteria for mobile hospitals	10
2.2 The location allocation problem and genetic algorithms	19
2.3 Mosul: city at the centre of a CHE	26
<b>3. Methodology</b>	<b>29</b>
3.1 SDSS overview	29
3.2 The setup phase	31
3.3 The safety phase	35
3.4 The genetic algorithm phase	36
3.5 Data	37
3.6 Tools	37
<b>4. Data preparation and model development</b>	<b>39</b>
4.1 Analysis of supplied candidate locations and generation of new locations	39
4.2 Calculating scores for evaluation	44
4.3 Implementing the safety phase	44
4.4 The genetic algorithm phase	46
4.5 User input and visualisation	55
<b>5. Results</b>	<b>57</b>
5.1 Quality and temporal variation	57
5.2 Comparison with previous models	70
<b>6. Discussion</b>	<b>78</b>
6.1 Quality and temporal variation	78
6.2 Comparison with previous models	80
<b>7 Conclusion, limitations and opportunities for future work</b>	<b>81</b>
7.1 Conclusions	81
7.2 Limitations and opportunities for future work	83
<b>8. References</b>	<b>85</b>

# 1. Introduction

## 1.1 Research context

Complex humanitarian emergencies (CHEs), which typically involve factors including large population movements, societal breakdown, political uncertainty, food shortages and severe disruption to infrastructure, are frequently catastrophic for public health (Toole and Waldman, 1997; Brennan and Nandy, 2001). Alongside the difficulty of fulfilling regular healthcare needs in a scenario where supply lines and facilities are likely to be compromised, medical professionals must respond to cases caused or exacerbated by the CHE itself, such as violent and psychological trauma, the spread of disease, malnutrition and sexual violence. Non-governmental organisations (NGOs), agencies, international bodies and military forces are often relied upon to fill this gap in public healthcare through deployment of temporary facilities such as mobile and field hospitals. These can play an important role in relieving the strain on what remains of the existing health network in the short to medium term.

Considering the threats posed by climate change, and continuing conflicts in Afghanistan, Syria and Iraq among other nations, healthcare provision in (complex) humanitarian emergencies continues to be pertinent. This research seeks to support the mitigation of the severity of public health catastrophes associated with such scenarios by taking a spatial approach to the location of mobile hospitals. To do this, it describes the creation of a robust spatial decision support system (SDSS) that can be used by decision makers who are not experts in geographical information systems (GIS) to rapidly identify sites in the CHE landscape where mobile hospitals can be most effective. The case study of Mosul, Iraq, which became a CHE during the battle for its liberation from the Islamic State of Iraq and the Levant (ISIS) in 2016 and 2017 was used to develop and test the decision tool.

### 1.1.1 Defining a mobile hospital

This work is centred on mobile hospitals so at the outset it is necessary to explicitly define what these constitute. This is something of a blind spot in existing literature with authors frequently conflating mobile and field hospitals due to their similar operating procedure and origins in military services. Further complexity results from the variation in type and materials that they use, implying that there is no typical field or mobile hospital (Manoochehry, 2018). However, emphasis on mobile hospital design by Bakowski (2016) suggests that differences in portability can be leveraged to make a nuanced separation.

For this project a mobile hospital was defined as a transportable healthcare unit that comes complete with the trucks and trailers necessary for relocation. A field hospital is also a transportable unit, situated in tents, containers or other structures, but does not include these vehicles. In practice ambiguity can be expected, with vehicles comprising some elements of a field hospital and mobile hospitals making use of exterior structures when in location, but the

distinction should suffice for the purposes of this research. The key point is that the design of a mobile hospital design is intended to give it greater maneuverability than its field hospital cousin. This definition is enhanced by an extensive discussion on the requirements and capabilities of mobile hospitals in section 2.1 of Chapter 2.

### **1.1.2 Location allocation**

The issue this work sought to address was essentially one of location allocation. Identifying the most suitable location for an amenity is a classic spatial problem, complicated by the priorities of different stakeholders engaged and invested in the outcome. Kahraman, Da Ruan and Doğan (2003) identify it as a decision problem involving multiple quantitative and qualitative criteria that is best addressed by heuristic techniques. Previous work in this area is extensive but solutions tend to be heavily dependent upon the context of a specific problem, meaning that applying a pre-existing SDSS to a new scenario is arduous. SDSSs are also often built on well-established GIS processes such as map algebra or network analysis, which are effective and come packaged with most GIS software but are not always the most efficient way of reaching an adequate solution. Traditional, highly tailored SDSSs are therefore not necessarily the best option for rapidly evolving CHEs where decision making speed can be critical to improving the situation in the field. Further to this, there is a shortage of literature attempting to find sites for healthcare facilities in CHEs, with facility location for effective supply lines - another important problem - receiving far more attention from geographers (for example Balcik and Beamon, 2008).

One of the ambitions of the project was to effectively utilise a method that avoided some of the shortcomings of these conventional SDSSs. With this in mind, heuristic algorithms are attractive because they do not need to try every combination of possibilities to deliver a result. These include genetic algorithms, which arrive at a solution through an iterative process inspired by natural selection (Larrañaga et al., 1996). They are able to take into account multiple criteria and with appropriate parameters are one of the fastest methods for resolving location allocation problems (Li and Yeh, 2005; Shariff, Moin and Omar, 2012). However, even with the existence of genetic algorithm frameworks in programming languages such as Python and R, they are not trivial to deploy; careful thought is required for specifying evaluation criteria and other parameters (Alp, Erkut and Drezner, 2003). Nevertheless, their speed and flexibility made them appealing for the project, providing they were applied in a manner that is understandable to users of the final application. Having illustrated why healthcare needs are so acute in CHEs and explained why genetic algorithms are an exciting possibility for location allocation of mobile hospitals, the project's formal objective and research questions will now be detailed.

## **1.2 Research objectives and sub-questions**

### **1.2.1 Overall objective**

The general objective of this research was to develop an SDSS that enabled the user to rapidly display indicative suitable location(s) for siting one or more mobile hospitals that met the needs of the greatest number of patients during and after a CHE.

This objective will be achieved by answering the following research questions and using these findings to carry out a practical implementation.

### **1.2.2 Research questions**

1. What are the requirements of a mobile hospital and in which order should these be prioritised?
2. How does a mobile hospital relate to the wider healthcare network, including Trauma Stabilisation Points and established hospitals?
3. Which algorithm can be used to solve location allocation problems quickly and effectively, factoring in dynamic variables such as movements of potential patients and disruption to transport networks?
4. To what extent do the results returned by the algorithm change as the CHE unfolds?
5. How can the algorithm be integrated with an SDSS to facilitate decision making for non-experts?
6. How do the results and performance of the algorithm compare to previous work in this field?

### 1.3 Scope

Despite formulating a clear research objective and research questions, as with any research there was a danger of being sidetracked and diverging from the project's central aim. As such it was useful to be explicit about areas related to the research that are out of scope. These were defined early in the project to mitigate any deviation from its core path and are listed here to manage the expectations of readers. They included:

1. Costs of operating a mobile hospital. Acknowledgement will be made of the agencies involved in delivering healthcare to the Mosul CHE, but this research is not intended to be an extensive discussion on the cost/benefits of deploying mobile hospitals to a CHE. Success of the mobile hospital was considered by determining the extent its location managed to meet the WMA's (2017) recommendations of saving the greatest number of lives while minimising morbidity.
2. The geopolitical situation of Mosul and Iraq more broadly. Although knowledge of the forces involved in the battle for Mosul was necessary to understand whether territory is friendly or hostile, this research used the Mosul CHE as an example context to develop an application. It does not offer significant commentary or analysis of the motivations of forces during the liberation.
3. Similarly, the effectiveness of the World Health Organisation's (WHO) response to the CHE was only considered where it had implications for the objectives of the project. The project does not offer detailed opinion or criticism of the WHO or other medical actors, although this is certainly a worthwhile activity in the general discourse about CHE healthcare. If the reader is interested, Fox, Stoddard, Harmer and Davidoff (2018) have carried out initial work in this area with a review of the healthcare operation including strengths and weaknesses.
4. Usability of the SDSS. While the SDSS sought to meet basic user requirements by clearly rendering a map and having some parameter options, carrying out an in-depth user needs analysis was out of scope for the project. It was stored on local files because the focus of the project is the algorithm underpinning the application. An avenue for future work could be a live version designed using results from a full user needs analysis.

## **1.4 Limitations at the outset**

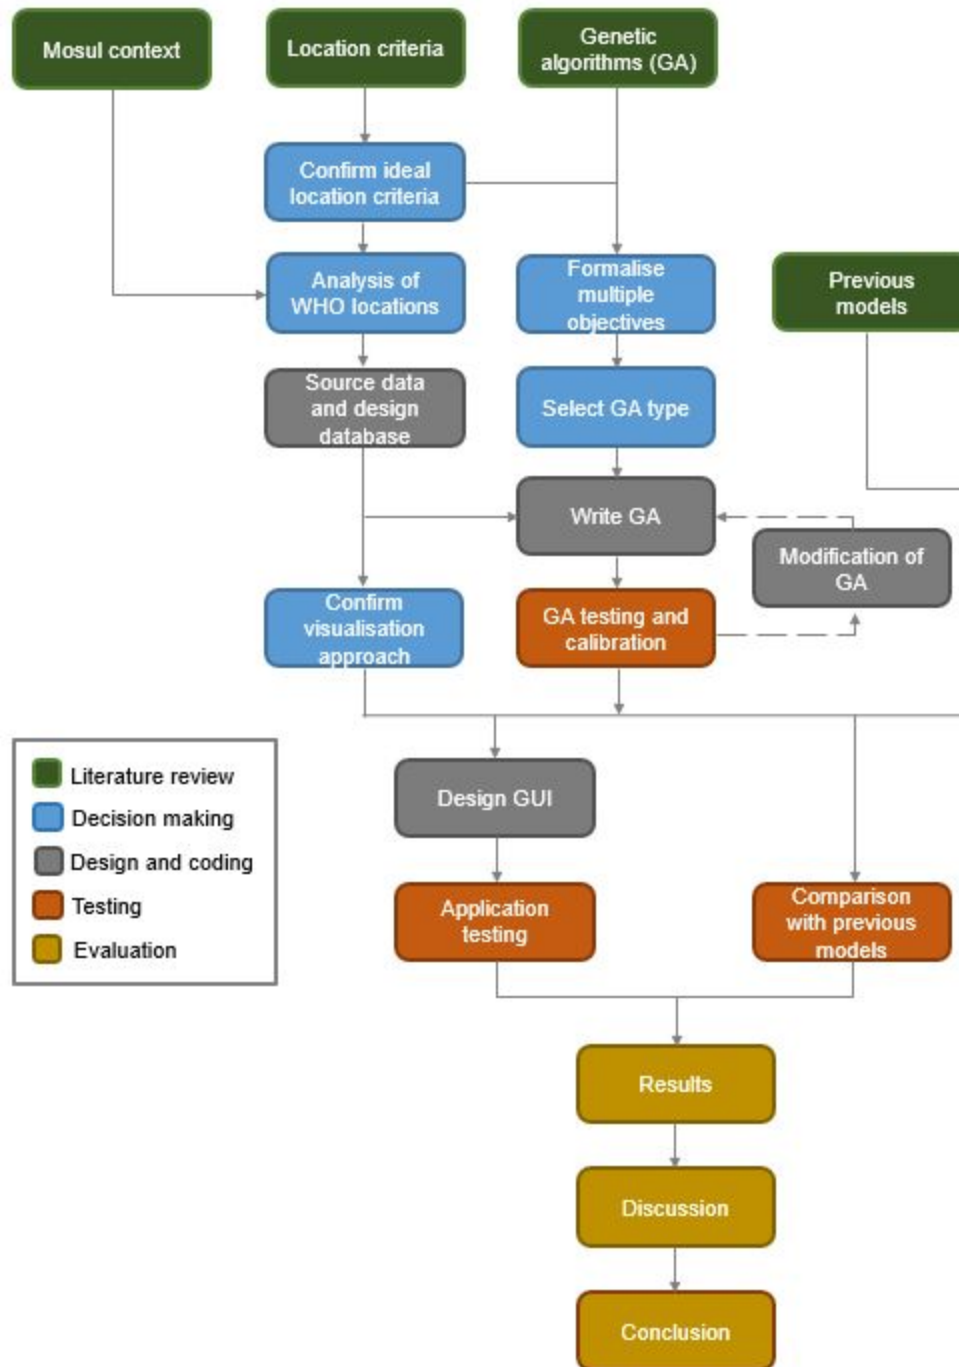
The final chapter includes a discussion of limitations of the SDSS itself and opportunities for future work, but some initial restrictions to the project were identified from the outset. These were:

1. While the SDSS sought to deliver locations with strong suitability for mobile hospitals, these were not guaranteed to be the absolute optimum locations. This is because the algorithm inside the SDSS utilised heuristic techniques to evaluate the multiple criteria for siting the facilities. Subchapter 2.2 elaborates on this and explains why these methods were appropriate.
2. The locations were intended to be advisory. A CHE is a dynamic environment with many potential dangers; although the SDSS was designed with this in mind, on-ground reconnaissance and expertise will be required to confirm the suitability of delivered results.
3. The SDSS was assessed using data from Mosul and dummy sets where appropriate. It is hoped that it can be repurposed for other CHEs but validation exercises will need to be conducted on the new data sets to confirm this.

## **1.5 Research framework**

To provide an overview of how the project proceeded, a graphical representation of the research framework is presented in Figure 1.1. The research followed a fairly straightforward flow chart structure, beginning with a literature review of three key elements, moving on to development and testing of the SDSS and finishing with results, discussion and conclusions. Choices regarding visualisation and testing were made in parallel to writing the algorithm the SDSS was built on, but were also informed by this aspect.





**Figure 1.1:** Research framework for the project.

## **1.6 Structure of this document**

This report describes the research project, following a structure complementing the initial steps shown in Figure 1.1. The current chapter has given the context to the work, explained the research gap it seeks to fill and detailed its objective, research questions and limitations. It is followed by a literature review in Chapter 2 that seeks to answer the theory focused research questions, providing material for developing the methodology underpinning the project as well as exploring the context of the Mosul conflict. Chapter 3 explains the methodology and the tools used, guiding the reader through the choices made when constructing the SDSS and deciding on the software with which to implement it. Chapter 4 details the tests carried out to establish the model's parameters and describes the model development through practical application on the Mosul case study. The results of experiments that aim to answer research questions four and six are presented in Chapter 5, along with commentary. The implications of these results are discussed in Chapter 6, which is then followed by a chapter containing conclusions, limitations and opportunities for future work. Finally, a full list of references is available in Chapter 8.

## **2. Literature review**


This chapter details the findings of literature reviews of three items: the location criteria that should be taken into account when siting a mobile hospital (2.1), methods for solving location allocation problems with particular focus on genetic algorithms (2.2) and the characteristics of the Mosul CHE (2.3).

### **2.1 Location criteria for mobile hospitals**

This section of the literature review intends to answer the first research question by drawing upon previous research and design material to identify the needs of a mobile hospital. By addressing this question, it should be possible to formulate a criteria list of requirements for a mobile hospital, some of which can be used as evaluation criteria in the SDSS. The review begins with a general discussion of mobile hospital characteristics before isolating areas of location criteria and considering each in turn.

#### **2.1.1 Mobile hospital structure and services**

Mobile hospitals fall within the Red Cross strategy of mobile health units, which are typically deployed when a population has no or little access to health care (ICRC, 2006). Within the health pyramid depicted in Figure 2.1 they are able to fulfil primary care needs relating to food security, water and sanitation, promotion of good health practice, rehabilitation and preventative activities. On top of this primary care function, some mobile hospitals are capable of offering more specialised services such as surgery, while others may emphasise emergency treatment over other types of healthcare provision (Bakowski, 2016). In a detailed review of healthcare solutions deployed under UN peacekeeping operations, Johnson (2015) argues convincingly for Level 2+ mobile hospitals, which deliver primary and secondary care while retaining capacity for enhancement with Level 3 features. These modular additions become significant when existing hospitals in the vicinity are compromised or overburdened. The standard Level 2 facility offers trauma care, general practice medical clinic and nursing services, while possible tertiary extensions include specialty care, extra beds and preventative medicine teams.



**Figure 2.1:** Health pyramid and care levels (adapted from ICRC, 2006, p. 9).

From a design perspective Bakowski (2016) also stresses the modularity of a mobile hospital, or the ease with which it can be packed into cubic containers. It is by definition an agile unit: the time costs associated with repositioning a mobile hospital are in general smaller than that of a field hospital. This is well illustrated by Médecins Sans Frontières' (MSF) Mobile Unit Surgical Trailer (MUST) which only takes three hours to set up or pack away (Hanlon, 2017). As a CHE unfolds it may be beneficial to take advantage of this feature and move a mobile hospital so it can avoid danger or better serve the population. However, if the primary difference between mobile and field hospitals is one of maneuverability, it is probable that useful information about mobile hospital location criteria can be gleaned from reviewing literature about field hospitals since they often are attempting to deliver similar levels of care.

Other examples of mobile hospital hardware include the Hospitrailers offered by Hospitainer, which make use of a modular shipping container design and are moved by articulated lorry (Hospitainer, n.d., a), and the mobile-home style integrated semi-trailer from the Medical Consultancy & Construction Group, also pulled by a lorry unit (MCC, n.d.). A mobile hospital is usually comprised of multiple trailers, as demonstrated by the five truck Hospitainer mobile surgical hospital that was dispatched to Mosul (Hospitainer, n.d., b). This multi-truck format is corroborated by Blackwell and Bosse (2007), who describe the MED-1 mobile hospital deployed to respond to Hurricane Katrina in the US. Two trailers were expanded on site to approximately 90m<sup>2</sup> using slide out pods and an awning area, and were surrounded by tents and trailers housing sleeping quarters, a command office, a canteen, and supplies. These existing mobile hospital designs have two immediate implications for the question addressed in this essay. Firstly, a reasonable amount of contiguous open space is required to host multiple articulated lorries in a layout that is optimal for the organisation of hospital facilities. Secondly, this area should be fairly level to guarantee safe working conditions for medics and patients inside the trailers, although there may be some capacity for adjustment using chocks.

Having discussed the characteristics of mobile hospitals, this review will now move on to factors to consider when locating them. With a relative shortage of literature on location allocation of mobile and field hospitals, work on siting permanent hospitals can provide a starting point for identifying these factors. As might be expected given the complex function a hospital plays in society, previous research has posited a wide range of criteria to consider, with categories including cost, demographics, market conditions or potential competition, management attitude, transportation, accessibility, availability of employees, availability of land for expansion, land ownership, existing infrastructure, air and noise pollution, proximity to services, and proximity to green space (Vahidnia, Alesheikh and Alimohammad, 2009; Senvar, Otay and Bolturk, 2016; Oppio et al., 2016; Wu, Lin and Chen, 2007). For mobile hospitals in a CHE some of these are clearly more relevant than others, and there are also issues unique to the CHE context that need to be recognised. The following subheadings present criteria that have been adapted accordingly, with discussion about why they ought to be taken into account in an SDSS.

### **2.1.2 Population**

For the expense of setting up a healthcare facility to be worthwhile, it needs to be situated in a place where it will be able to treat the greatest number of people, and where there is a gap in the existing health infrastructure serving this population. Leaving aside the latter point for the moment, the crudest way to satisfy the former criteria would be to locate the mobile hospital in a place where the greatest number of people can access it. Indeed, population counts and densities are used in tools to position permanent hospitals (see Vahidnia et al., 2009), implying that they should be considered in the SDSS. However, this data does not necessarily accurately reflect differences in the level of need across the population for such a facility. Senvar et al. (2016) take a more nuanced approach, incorporating age profiles and socio-economic status to determine areas with the highest expected demand. These locations may well not correspond exactly with population distribution and instead favour higher densities of vulnerable individuals such as the elderly or young children.

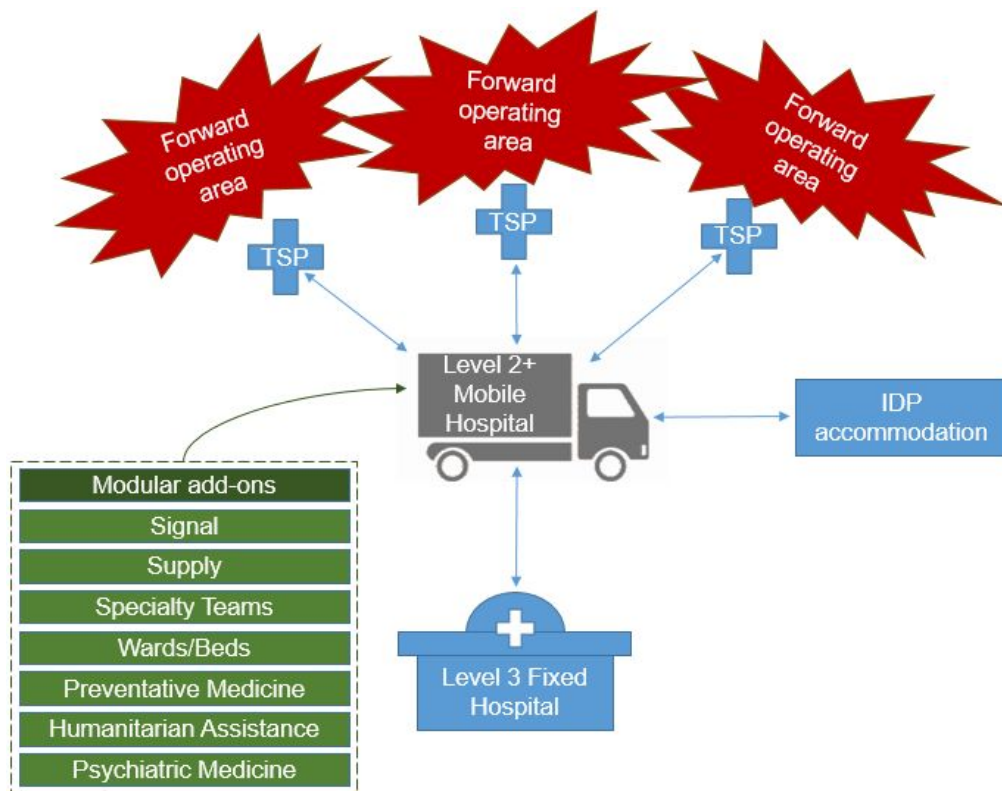
In a CHE, demand distribution becomes harder to anticipate because healthcare issues caused or exacerbated by the emergency itself are added to the regular needs of the population. Primary care needs relating to food security and sanitation, which tend not to impact the majority of the population in a non-emergency scenario, are suddenly of significant concern. Demand is further heightened by traumatic injury rates, which increase substantially among civilians in modern conflicts (Hinsley et al., 2005). Areas of high injury concentration should be considered when positioning the mobile hospital, although other services in the healthcare network like Trauma Stabilisation Points (TSPs) may meet the immediate needs of these patients.

Another variable related to population affected by a CHE is people migrating en masse from conflict zones in pursuit of safety and stability. The act of moving itself along with the mixed quality of destination accommodation can act as catalysts for outbreaks of infectious diseases (Toole and Waldman, 1997; Raad et al. 2018). The adequacy of static census data on the

population prior to the CHE is consequently limited for the SDSS as distribution and healthcare demands can change quickly as the emergency unfolds. Reliable information on the movements of IDPs is a big asset for choosing candidate mobile hospital locations because it enables the tool to ascertain where the greatest number of people are located and also where this intersects with those who are likely to most urgently require healthcare.

### **2.1.3 Healthcare network**

As alluded to by the health pyramid and care levels diagram in Figure 2.1, healthcare facilities do not exist in isolation and instead interact together to deliver different types of services to the population. This relationship forms what the WHO terms 'integrated health services', which include holistic policy-making and management, vertical integration of different levels of service, multi-purpose service delivery points and overarching strategies of intervention for certain population groups (WHO, 2008). Spatially, integrated health services can be viewed as a network, with a few large university hospitals complimented by many more health centres and health posts. In an emergency scenario it is important that any temporary facilities like TSPs or mobile hospitals are sited in locations facilitate harmonious interaction with the existing network (Pérez, 2015). A mobile hospital, which can provide primary care and some specialised secondary or tertiary services, has the potential to act as a middle node in this network, treating patients referred from smaller facilities and passing others on to permanent hospitals when necessary. Figure 2.2 visualises these relationships and some of the services that might be added to a mobile hospital, using a model originally posited by Johnson (2015) while advocating hybrid Level 2+ mobile hospitals.



**Figure 2.2:** Mobile hospital within a CHE healthcare network (adapted from Johnson, 2015) showing its relationships with TSPs, internally displaced people (IDPs) and existing hospitals.

In a CHE, one of the main functions of the integrated health services network - or what remains of it - is to provide a smooth medical referral pathway so that victims of traumatic injuries successfully reach a facility capable of providing the required treatment. The 'golden hour' principle, whereby the chance of a patient surviving decreases significantly if they are not transferred to a care facility within an hour of sustaining an injury, is one possible metric that can be considered here. Its popularity is challenged by Rogers, Rittenhouse and Gross (2015), who argue that it is overly simplistic, leading drivers to take unnecessary risks when attempting to deliver patients within this time window, heightening the possibility of additional problems. Travel time from the area of fighting to a care facility is still paramount, but perhaps a conservative average speed should be used when calculating the 'golden hour' from a candidate location to TSPs. The possibility of roads being in poor repair or blocked adds further weight to this point.

As argued by Johnson (2015), mobile hospitals should also be able to refer patients to facilities providing higher care levels, making existing hospitals a valuable inclusion in the SDSS. This was an explicit intention of the WHO's coordinated trauma response in Mosul (Fox et al., 2018).

The distance factor is slightly more complicated in this instance because some of the services offered by the mobile hospital will be the same as those in the permanent facility, so close proximity will not be the best spatial arrangement of these resources for an efficient integrated health service network. A detailed understanding of the types of services available at each node in the network and their capacity would therefore be useful for determining the optimal placement of a mobile hospital.

#### **2.1.4 Accessibility and transport network**

The efficiency of integrated health services, especially the effectiveness of the medical referral pathway, is significantly influenced by the quality of the transportation network, as alluded to by the connection arrows in Figure 2.2. Vahidnia et al. (2009) take this into account in their SDSS by favouring sites near arterial streets. In a CHE where public transport is likely to be severely disrupted, patients make use of cars, walking, carts and in some cases improvised ambulances to reach treatment facilities. However, CHE conditions can compromise the existing road network, temporarily through obstructions like security roadblocks, ongoing fighting, aerial bombardment and broken down vehicles, or more long-term due to key links such as bridges being destroyed. This is alluded to by Salihiy (2019) when describing the horrendous conditions endured by civilians leaving ISIS controlled areas in Mosul, frequently pursued by militants and then put through highly problematic security screenings.

A mobile hospital ought to be located somewhere as accessible as possible, firstly so that trucks containing the facility can actually arrive there, and secondly so that patients are able to reach the site. At a minimum the SDSS should include criteria on distance from roads, but to better reflect connectivity and how the CHE effects road segments over time, a more sophisticated network approach ought to be taken to find sites that stand the best chance of staying connected to potential patients, TSPs and other hospitals.

#### **2.1.5 Safety concerns**

In non-emergency situations safety concerns for hospital site selection include the health risks posed by air pollution (Vahidnia, Alesheikh and Alimohammadi, 2009) and hydrological instability (Oppio et al., 2016). These could also be acknowledged in a CHE SDSS, but clearly there are more severe dangers to the wellbeing of staff and patients in an emergency context. A survey of disaster medical assistance teams furthers understanding here, calling for the security of the humanitarian community to be prioritised (Aitken et al., 2009), especially given the trend of medical workers being viewed as legitimate targets in armed conflicts (Bricknell and MacCormack, 2005).

The relative maneuverability of a mobile hospital means that it can, if necessary, relocate if hostile forces advance, but this capability to retreat is not a watertight safety protocol in a fast-moving warzone situation. Training and briefings, coordination with friendly military forces and a clear hierarchy of responsibility for security are all effective measures for mitigating



security risks (Aitken et al., 2009), but spatial factors should not be neglected. Some authors recommend designating safe zones and corridors for locating humanitarian facilities, with the caveat that these rely on agreement from all parties in the conflict and that they should be developed in collaboration rather than imposed on an area by an external force (Gilbert and Rusch, 2017). Safe zones are appealing for the mobile hospital location allocation problem but there is no guarantee that all CHEs will have such clearly defined areas, and there is still a danger that some actors may violate their sanctity.

Perhaps the most straightforward approach is to ensure that there is a spatial buffer between the facility and the frontline. However, territory ownership is frequently fragmented and ambiguous in contemporary urban warfare, and in some CHEs safety at the hospital may be compromised by long range artillery and airstrikes. A more nuanced approach could incorporate Schutte's (2017) finding that indiscriminate violence in civil wars is most likely to occur close to the power centre of military forces. In addition, UN guidance on military and humanitarian interaction explicitly stresses that there should be a clear distinction between the two actors, which prohibits co-location (OCHA, 2017). This adds weight to the argument that for safety purposes mobile hospitals should not be located near military bases and infrastructure.

### **2.1.6 Land characteristics**

Research on locating permanent hospitals understandably devotes significant attention to land ownership, value and current use (Oppio et al., 2016). In a CHE the social structures previously governing these factors may have been superseded by military law, which in theory informs decisions about when it is appropriate to commandeer property and the types of activities that should be permitted at these locations. Occupation law does exist in a codified form but involves a complicated legal framework and can be difficult to apply in practice (Scheffer, 2003). Under US policy, for example, seized religious buildings should only be used for medical needs (Santerre, 1989).

In their model for locating field hospitals, Vafaei and Oztaysi (2014) avoid contention over land ownership by only allowing public parks to be selected as candidate locations. This is pragmatic but arguably too restrictive. Drawing upon work that considers the location of other amenities during an emergency situation, Cetinkaya et al. (2016) present a helpful criteria pool for siting a refugee camp. Public land is deemed preferable, woodland is excluded to minimise illicit tree felling, and windy areas are avoided to prevent damage to the tents and containers that form the camps. In addition, the authors favour accessible plane areas with a slope no greater than 7%. For CHEs occurring in diverse landscapes where there are big changes in elevation these specifications may be well worth repurposing because they are broadly applicable to mobile hospitals.

However, without meaning to underplay the frictions and challenges that can arise between occupying forces, aid agencies and civilians in situ, the emphasis of this project is not on the legality of hospital sites. It is reasonable to expect that negotiation with stakeholders on the

ground will be required before deploying a mobile hospital at any indicative optimal location provided by the SDSS. Fortunately, the self-contained designs of mobile hospitals lend themselves to deployment in open spaces, an arrangement that is perhaps less likely to cause controversy than setting up shop in an existing structure.

### **2.1.7 Utility networks**

Water and electricity are needed to deliver effective healthcare, and in a CHE it can be expected that the infrastructure providing these utilities will be disrupted or perhaps damaged beyond repair (Brennan and Nandy, 2001). Mobile hospital design does take this into account, with MSF's MUST unit able to operate fully autonomously for a week before needing to resupply with fuel and water (Hanlon, 2017). The UN Level 2 Mobile Hospital designed by Hospitainer also boasts an energy supply and water purification system within its specifications, although it is unclear how long it can work without any sort of replenishment (Hospitainer, n.d., c).

Although the ability to carry out effective work off-grid is a useful trait of mobile hospitals, unsurprisingly academic literature on field and mobile hospitals indicates that it is preferable to tap into utility networks if available. A discussion of a field hospital established to provide support following an earthquake that struck Chile in 2010 notes that a sports complex was an excellent location partly because the perimeter lighting circuit provided a good electrical connection and water systems were also available (Pérez, 2015). No significant issues with these supplies are mentioned but a generator was able to provide around 12 hours of power in the event of a blackout. Elsewhere, a field hospital deployed to Haiti in 2010 also came equipped with generators to guarantee a stable electricity supply (Kreiss et al., 2010) as did the Carolina MED-1 mobile hospital sent to support victims of Hurricane Katrina (Blackwell and Bosse, 2007).

The various designs of mobile hospitals and evidence from the case studies of field and mobile hospitals indicates that while connecting to reliable electricity and water supplies is preferable, generators and water tanks mean that it is not an essential requirement. For the SDSS, data layers recognising these networks would be useful to include but do not need to be given a high weighting when calculating candidate locations.

### **2.1.8 Final criteria overview**

Having discussed various criteria that could be considered when deciding on the location of a mobile hospital, Table 2.1 lists those that might feature in the SDSS. These are separated by category as in the literature review. Notes on importance are also included as a first step towards selecting the criteria for the final tool.

**Table 2.1:** Criteria for possible inclusion in the SDSS.

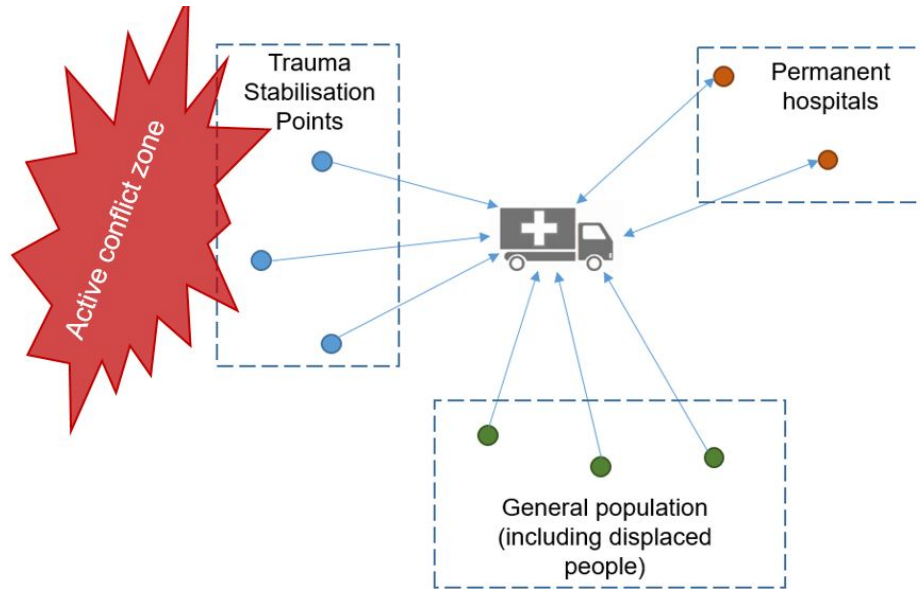
<b>Category</b>	<b>Criteria</b>	<b>Notes</b>
Population	Density	Very important. Will need the capacity to update these as CHE unfolds.
	Location of IDPs	
	Areas where population is disproportionately impacted by violence (increased traumatic injury rates).	
Healthcare network	Proximity to hospitals	Need to be within acceptable travel time so patients can be referred to more sophisticated facilities.
	Proximity to Trauma Stabilisation Points	Very important. Need to be within the 'golden hour' window of travel time from multiple TSPs so trauma victims have best chance of survival.
Accessibility	Proximity to major road	Very important for arrival of mobile hospital and incoming patients.
	Availability of alternative routes if main road is compromised	Not crucial but would be sensible to have at least one alternative access point.
Safety	Distance from conflict zone	Very important for there to be a buffer distance between mobile hospital and conflict zone.
	Natural hazard risk	Not a priority in the Mosul context but in some CHEs may be important, particularly considerations such as flood and landslide risk.
	Proximity to military installations.	Not critical but would be useful to mitigate conflation with friendly military forces.
Land	Size of contiguous space	Needs to be big enough to

characteristics		accommodate layout of trucks and tents.
	Open space	Should not be forested but will be possible to work around some features (individual trees).
	Ownership	Not essential but public ownership should be favoured over private.
	Minimal slope	Not a major issue in Mosul, which mostly lies on a plain.
Utility networks	Water supply	If regular supplies are available then these ought to be included. Preference given to proximity to the networks.
	Electricity supply	

**2.2 The location allocation problem and genetic algorithms**

The choice of location for an amenity has a very influential bearing on it achieving its objectives, whether these are financial or philanthropic. The resources required for set-up and the possible prohibitive costs of relocation add further weight to this initial location decision. Site selection is further complicated by the requirements of the amenity and input from associated stakeholders, some of whom may have conflicting demands. Systematic approaches have been developed to aid decision makers by returning a suggested location, but there is no definitive solution because each location allocation problem comes with a distinct context.

As explored in the previous sub-chapter, mobile hospitals have a whole host of location requirements, with criteria relating to safety, accessibility and proximity to utility networks. Aside from safety, perhaps the most important factor is connectivity to the population who need its services and to other components of healthcare infrastructure. As such, the location problem can be visualised in terms of the network illustrated in Figure 2.3. The mobile hospital must be positioned at a point safely outside the active conflict zone where it can accommodate trauma victims, attempt to satisfy the healthcare needs of the general population and make referrals to permanent hospitals. Travel time to all these nodes should be minimised, although in practice it is expected that certain node types should be favoured.



**Figure 2.3:** Mobile hospital within a simplified version of a CHE healthcare network.

This section of the literature review will first discuss network-centred approaches to location allocation, focusing on the  $p$ -median problem. It will then offer a detailed consideration of the strengths and weaknesses of applying genetic algorithms to solve this problem and location allocation in other contexts. To make sure the review is wide ranging, the focus will not exclusively be on work determining sites for healthcare facilities.

### 2.2.1 Defining the $p$ -median problem

Academic work addressing location allocation from a network perspective has its foundations in maths and economics. One of the earliest formal attempts at solving a location allocation problem is that of Friedrich and Weber (1929) who calculated the optimal position of a warehouse to minimise the travel distance between it and multiple customers. Authors like Cooper (1963) continued this research by presenting exact extremal equations and an approximate matrix solution for solving the problem of efficiently arranging multiple sources (distribution points) among fixed destination points. A year later, Hakimi (1964) made a highly influential contribution through his formulation of the  $p$ -median problem.

The  $p$ -median problem is easy to comprehend but less straightforward to solve. In an environment with  $n$  demand points, the problem requires computation of the combination of  $p$  supply points with minimal total distance of each demand point to its closest supply point (Jackson, Rouskas and Stallmann, 2007). Importantly for this project, the model recognises the deterministic nature of travel times along vertices and demand for services at nodes. As

demonstrated in Figure 2.3, locating mobile hospitals in a CHE can also be visualised as a  $p$ -median problem of sorts.

Much work has built upon Hakimi's original premise, but of particular interest is Serra and Marianov (1998) who focused on demand uncertainty in the problem. They considered fire station locations in Barcelona, a city where there is considerable temporal variation throughout the day in the population distribution due to commuting and tourism patterns. They addressed the problem by reflecting temporal changes through different scenarios and by defining two objectives (minimise the maximum average travel time across scenarios, and minimise the maximum regret) to work against. Evidently it is possible to use the  $p$ -median model in contexts where there are frequent changes to the networks and multiple objectives for the supply points.

Although the  $p$ -median problem is well established, even with contemporary computing power finding an exact solution is usually impractical. Soft computing methods such as heuristics are capable of delivering an approximate solution within a reasonable time threshold, but choosing the appropriate method is challenging because a range of heuristics and metaheuristics are available (Mladenović, Brimberg, Hansen and Moreno-Pérez, 2007). The speed of returned results is an explicit priority of this project, making the genetic algorithm an attractive metaheuristic to explore in further detail because its efficiency has been well demonstrated (Alp et al., 2003).

### **2.2.2 Genetic algorithms**

Genetic algorithms use the mechanics of natural selection to identify the 'fittest' solution (Holland, 1992). Hosage and Goodchild (1986) were the first to apply them to the  $p$ -median problem, but they have also been used in a considerable number of other fields, including land allocation (Schwaab et al., 2018), image processing, gaming and task prioritisation (Kumar, Husian, Upreti and Gupta, 2010). Larrañaga et al. (1996) provide a concise overview of how they operate. From an initial population of candidates, parents are selected. Combinations (crossovers) of these parents result in children who inherit characteristics of both of their parents. There may also be capacity for mutation at this stage depending on the specific algorithm. The children are added into the population, all of which is tested for fitness, and the least fit members are removed to ensure the total number remains the same. The overall quality of the population should therefore increase as crossovers are carried out. This process continues until the best solution is identified, the algorithm hits a predefined maximum number of generations, or a member of the population satisfies some predetermined condition. Table 2.2 uses pseudocode to outline this sequence.

**Table 2.2:** Genetic algorithm pseudocode (adapted from Toutouh, Rossit and Nesmachnow, 2019)

Pseudocode	Function
1: $t \leftarrow 0$	Generation counter
2: initialise( $P(0)$ )	Population initialisation
3: While not stopping_criterion do	While statement. Process repeats until stopping_criterion is fulfilled (usually fitness of population or number of generations).
4:     evaluate( $P(t)$ )	Population evaluation
5:     parents $\leftarrow$ selection( $P(t)$ )	Select parents for crossover, different sampling methods available to do this.
6:     offspring $\leftarrow$ evolutionary operators(parents)	Produce two offspring using evolutionary operator to decide which characteristics will be taken from parents.
7:     (optional) mutate offspring	Offspring mutated using an external operator.
8: $P(t+1) \leftarrow$ replacement(offspring, $P(t)$ )	Insert offspring into population, remove, replace current population with new one.
9: $t = t + 1$	1 added to generation counter
10: end while	
11: return best solution found	

It is important to note that while genetic algorithms typically follow the above process, there is considerable opportunity for modification in each research case. Toolkits are available in multiple programming languages, but users are required to make design decisions about factors like the size of the population in the generation, how the crossover operator works, mutation opportunities and the termination criteria (Alp et al., 2003). That said, if well thought through, genetic algorithms can be a powerful and effective way of finding a near-optimal location, as will be discussed next.

### 2.2.3 Usefulness of genetic algorithms for location allocation

Various authors have conducted thorough tests which add weight to arguments in favour of genetic algorithms for location allocation. Houck, Joines and Kay (1996) compared a genetic algorithm with random restart and two-opt switching methods. For small test problems the results of all three were comparable, but when addressing larger problems the genetic algorithm was able to deliver better solutions in less time. This was due to its ability to combine characteristics of two attempted solutions to devise an improved solution instead of a blind search approach, which requires analysis of all possible solutions.

Li and Yeh (2005) went further by comparing a genetic algorithm with neighbourhood search and simulated annealing to solve a real-world problem of allocating hospitals in Hong Kong. A GIS was used to calculate population and transportation constraints influencing the fitness of potential facility locations. The genetic algorithm provided superior results to the other two techniques, with computation time only being an impressive 29.4% of simulated annealing. The research also demonstrated that a genetic algorithm can be combined with multicriteria evaluation techniques by using a linear weighted equation in the algorithm. Elsewhere, simulated annealing appeared to provide a slightly better solution than a genetic algorithm to the problem of devising service areas for police forces (Duarte, Henriques and Ribeiro, 2019). However, no commentary was provided on differences in execution time or computing power required for each technique.

Genetic algorithms can also be applied to calculate maximal coverage, a common sub-category of the location allocation problem. A study by Shariff, Moin and Omar (2012) considers the issue of healthcare facility allocation, viewing the problem through a network prism with demand nodes and facility nodes connected by vectors of differing lengths. A further interesting takeaway from this paper is the decision to use the benchmark problems devised by Haghani (1996) to test the algorithm before applying it to the actual case study. The merits for future research of testing against an established set of problems are clear, although these trials should not be assumed to be finite.

While the above discussion identifies some of the advantages and successful application of genetic algorithms, their weaknesses should also be acknowledged. As already mentioned, genetic algorithms are heuristics, meaning that they do not return the same final result to the problem every time they are run and that the result they do deliver is not guaranteed to be optimal. For location allocation problems attempting to satisfy multiple quantitative and qualitative criteria, some of which may be conflicting, this is a tolerable drawback because finding the ultimate optimal solution may not be practical or worth the level of rigor required.

Another disadvantage of genetic algorithms is their sensitivity to the size and diversity of the initial population. Diaz-Gomez and Hougen (2007) note the difficulty in striking a balance between a small population, which can deliver poor solutions, and the amount of computing power required to calculate a large population. They propose alternative ways of measuring



population diversity, favouring the centre of mass as a metric. The final area of weakness may seem obvious but is still worth recognising, especially in a contemporary world where even complex problems can often be solved by out-of-the-box software solutions. The strength of a genetic algorithm is dependent on design decisions, especially regarding the crossover operator, mutation probability and the stop criteria. Therefore the time needed to write, test and calibrate the algorithm should not be underestimated. This can also risk over specialising: an algorithm designed for one location allocation problem may be completely inappropriate for another, even if the context is similar.

#### **2.2.4 Recent developments**

With genetic algorithms shown to offer a viable solution to the  $p$ -median problem, recent work has focused on incorporation of external factors influencing demand, supply, travel time and installation costs. Huang and Wen (2014) worked out the optimal spatial distribution of which 7-Eleven stores in Taipei, Taiwan defibrillators should be placed in, incorporating spatial and temporal data on cardiac arrests. They also added a stirring operator to mitigate premature convergence. Sharifi Noorian, Psyllidis and Bozzon (2018) returned to the issue of travel cost changing over time, adding uncertainty to the  $p$ -median model. Working in the context of fish markets as suppliers and restaurants as demand areas, they devised three temporal scenarios over the course of a regular Monday which each had different levels of traffic disruption. The fitness function aspect of the genetic algorithm assessed each demand area to see if it was allocated the closest market at the specific temporal scenario with minimal travel cost distance. This resulted in a new configuration of fish markets and demand areas theoretically more time-efficient for the restaurants. Although the authors recognise limitations with regard to data quality, this work demonstrates that genetic algorithms can successfully accommodate external factors.

Similarly, Toutouh, Rossit and Nesmachnow (2019) compared multiobjective evolutionary algorithms for locating garbage accumulation points and the number and type of bins at each point. The intention was to maximise waste collected, minimise expenditure on bins and minimise the distance between rubbish creators and bins. Non-dominated Sorting Genetic Algorithm, version II (NSGA-II) and Strength Pareto Evolutionary Algorithm, version II were adapted to address the issue. Analysis suggested that the former was more successful, delivering superior solutions. Zeng et al. (2019) also used NSGA-II to pick a suitable combination of electric charging points to expand infrastructure for electric vehicles, with the multiple objectives of minimising costs and number of uncompleted trips. This was preceded by a cluster algorithm that returned candidate locations, something that is worth emphasising as clearly such algorithms cannot operate without a list of potential locations.

## 2.2.5 Tools for creating genetic algorithms

Alp et al. (2003) were keen to stress that writing a genetic algorithm is not a trivial matter. This observation remains true even with the advances in the field and computing more broadly in the fifteen years since it was made. As a precursor to writing an algorithm, a useful exercise is looking at the programming languages and software used in the literature because this could help steer the process. Table 2.3 presents the results of this process with articles ordered chronologically. Note that this list is far from exhaustive, it is merely intended to illustrate the different options available for constructing a genetic algorithm.

**Table 2.3:** Languages and software used to write genetic algorithms.

Author(s)	Algorithm language and software details.	Rationale
Alp et al. (2003).	Written in C++	Unspecified
Li and Yeh (2005).	Written in Visual Basic. GeneHunter package used for evolutionary aspect.	GeneHunter can call genetic algorithm functions flexibly, also accommodates common programming languages.
Shariff, Moin and Omar (2012).	Unspecified.	Unspecified.
Huang and Wen (2014).	Written in Python. QGIS used for inputs and outputs.	Considered a good prospect for further simulation experiments.
Sharifi Noorian, Psyllidis and Bozzon (2018).	Written in Java 8.	Unspecified.
Toutouh, Rossit and Nesmachnow (2019).	Written in Python using the DEAP library.	Unspecified.
Zeng et al. (2019).	Written in Python using the DEAP library.	Unspecified.

## 2.2.6 Points to consider for the SDSS

The evidence suggests that a genetic algorithm can be an efficient and appropriate method for addressing location allocation problems, particularly when there are multiple criteria involved. The availability of existing frameworks such as Distributed Evolutionary Algorithm in Python (DEAP) furthers the appeal of using a genetic algorithm to power the SDSS. However, caution needs to be exercised during the design and testing process: the various design options may lead to the SDSS being over-specialised for the context of the Mosul CHE.

## **2.3 Mosul: city at the centre of a CHE**

The SDSS that results from this project is intended to be sufficiently generic that it can be applied to future CHEs with a minimum amount of modification. However, data is still required to write and test the underlying algorithm, so it is sensible to draw upon a previous CHE that was relatively well documented. The city of Mosul in Ninewa, Iraq, which faced a public health emergency during its liberation from ISIS, has been chosen because it is a recent CHE example and was under close observation from the WHO, meaning reasonable quality data is easily available. This section of the literature review explains the case study with special attention given to healthcare provision and factors that might influence mobile hospital location.

### **2.3.1 Origins of the CHE**

Reclamation of territory in Mosul from ISIS by a coalition that included Iraqi state and Kurdish soldiers took place in three phases between October 2016 and July 2017, beginning in the east of the city. A military analysis of the fighting stresses how difficult it was for forces to gain ground because ISIS defenders used the city's dense urban environment to control the rhythm and conditions of combat. Approximately half of the nine month battle was spent on the third phase of the assault, which aimed to capture west Mosul. During this stage ISIS inflicted 7,000 coalition casualties, a significant proportion (85%) of the total (Arnold and Fiore, 2019).

For civilians, circumstances during the fighting could certainly be categorised as a CHE, a point corroborated by a survey of 1,202 households across the city conducted by Lafta, Al-Nuaimi and Burnham (2018) in 2017. Matching the phases of the assault, injury, loss of life and damage to buildings varied between the east and west of the city, with the latter disproportionately affected by airstrikes. The mortality rate increased from 0.58 deaths per 1,000 person-months under ISIS control to 6.29 in east Mosul during the fighting, and 0.64 to 15.54 in the west of the city. Over half of households surveyed experienced building damage, while 86% in west Mosul reported no electricity. Almost all households relied on water drawn from wells due to pipelines being compromised.

### **2.3.2 Population movements**

Internal migration triggered by the ISIS occupation means there is some ambiguity about the size of the civilian population prior to the battle, but estimates suggest almost 1.4 million inhabitants (UN-Habitat, 2016). A dedicated displacement tracking matrix (DTM) portal reveals that in excess of 830,000 individuals were displaced from the city during the fighting, 88% of whom came from west Mosul (IOM, 2017). 95% of internally displaced persons (IDPs) did not leave the wider Ninewa governorate, with 77% staying in the district of Mosul. As of June 2017 over 90% of IDPs were categorised as living in emergency sites, camps or unknown types of shelter (IOM, 2017), precarious accommodation types with healthcare, sanitation and utility

challenges. Alongside these population movements, deciding where to direct resources including healthcare was complicated by a shifting frontline and difficulties defining boundaries inside the city (Arnold and Fiore, 2019).

### **2.3.3 Healthcare provision**

Under ISIS, Mosul's healthcare network was already facing significant challenges to delivering effective care at all levels. An interview of twenty healthcare employees who worked under the regime exposes the severity of conditions, with themes including random deaths and kidnappings, psychological distress and poor availability of medicine and equipment (Michlig, Lafta, Al-Nuaimi and Burnham, 2019). Although some journalists have expressed surprise at the extent ISIS strongholds such as Mosul functioned bureaucratically under such a grim regime (Callimachi and Mills, 2018), evidently the healthcare needs of the city's population were not being met even before any fighting began.

When planning the invasion the Iraqi military recognised this and the additional healthcare burden that would manifest as a result of an aggressive intervention but acknowledged that they had insufficient medical resources to treat civilians. Healthcare was therefore coordinated by the WHO, which contracted two NGOs (NYC Medics and Samaritan's Purse) and private supplier Aspen Medical to deploy medical teams (Beaubien, 2018). Field hospitals were established at three satellite towns to the south and east of Mosul: Bartella (January 2017), Adhba (March 2017) and Hamam al Alil (April 2017). Security concerns were a major motivation for choosing these sites, and the Bartella and Adhba locations were both criticised due to their poor accessibility (Spiegel et al., 2018). The Samaritan's Purse field hospital at Bartalla was also reprimanded for the militarisation of its appearance and the vigorous security checks demanded of patients (Fox et al., 2018).

To complete the medical referral pathway network, TSPs were set up within ten minutes travel of the frontline to process incoming trauma patients before dispatching them to field hospitals (Fox et al., 2018). Amouri and Reed (2018) point out that these transfers often resulted in patients dying from manageable and preventable causes because they frequently took place in regular vehicles lacking medical equipment and could take between 30 and 90 minutes. Despite this issue it is estimated that the medical referral pathway system saved 1,500 - 1,800 lives (Spiegel et al., 2018). However, it also reiterates the importance of careful location selection for mobile hospitals and consideration of how they fit into the wider network of medical facilities and transportation.

### **2.3.4 Points to consider for the SDSS**

To summarise, there are a few key points from the Mosul CHE that will need to be reflected in the SDSS. Firstly, fighting in the Ninewa region during the battle was spatially concentrated on the city itself, and for much of the time on the western districts of the city. With hindsight, it can be seen that the eastern side of the city remained relatively stable once it had been captured by

coalition forces but when the battle was in progress there would have been no guarantee of this being maintained. Secondly, the case study reinforces arguments that civilian healthcare needs during a CHE are complex. The civilian population who remained in the city suffered traumatic injuries while lacking general healthcare provision. IDPs in the hinterland may have escaped the worst of the physical violence but were vulnerable to healthcare challenges caused by difficult journeys and inadequate accommodation. Finally, although the work of the field hospitals was undoubtedly admirable, shortcomings regarding accessibility and the safe transfer of patients from TSPs should be recognised.

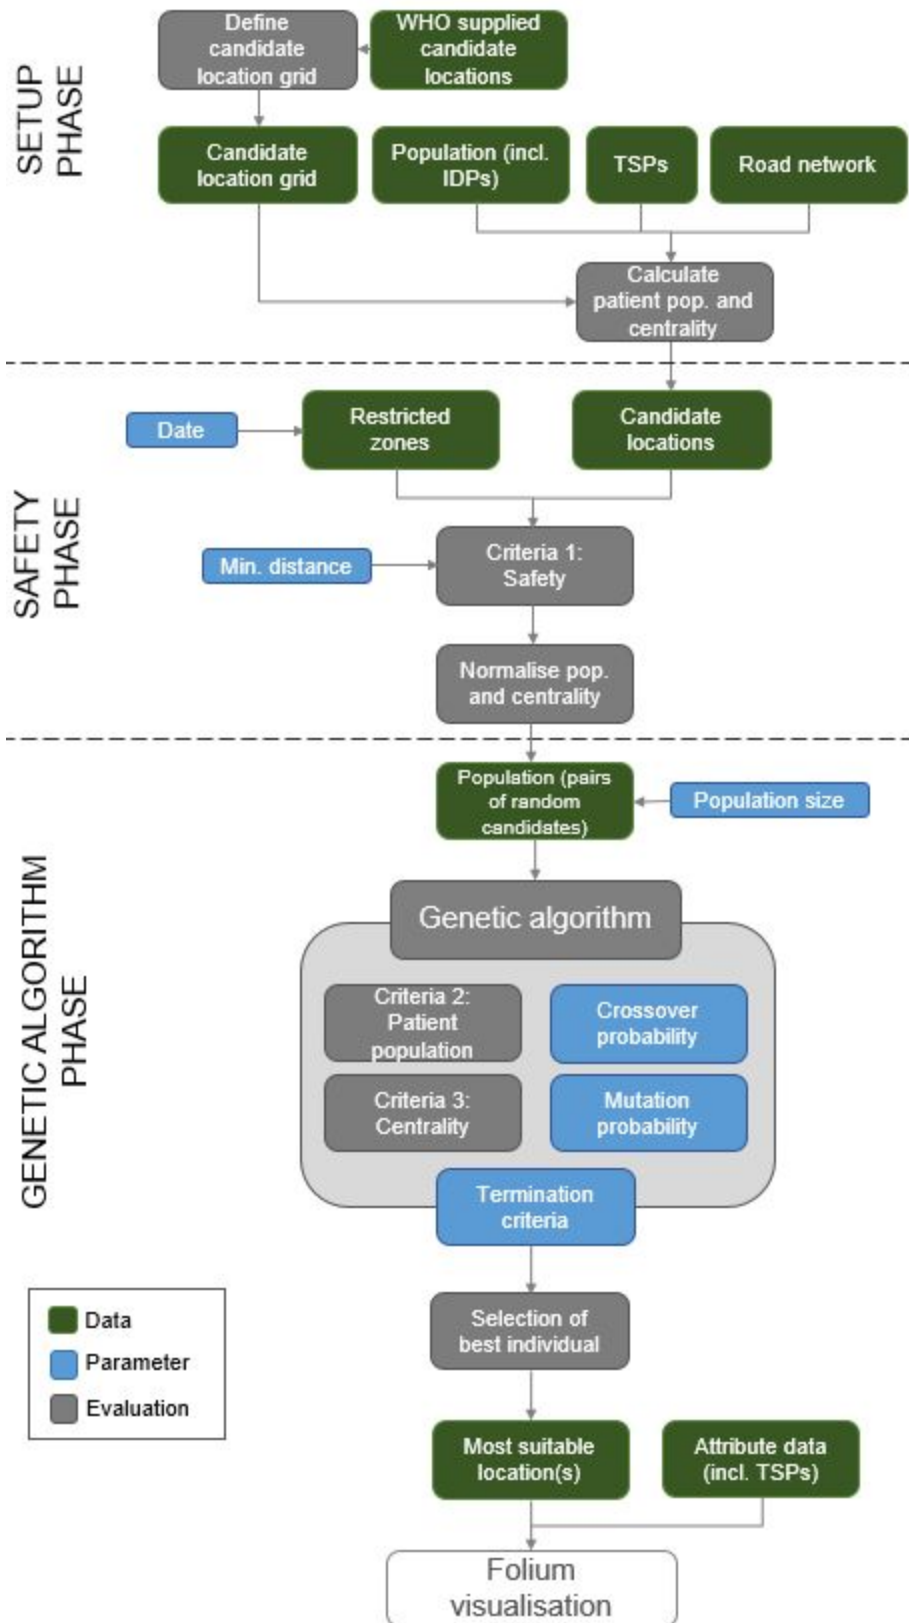
### **3. Methodology**

Informed by the findings of the literature review, this chapter details the methods, data and tools used to create an SDSS that can deliver suitable locations for mobile hospitals. It begins with an overview of the different components that form the SDSS before explaining how each of these were derived and the rationale behind the selection of tools and data. The following chapter details further development of the SDSS when it was applied to the Mosul case study, particularly the finalisation of various parameters that influence the behaviour of the genetic algorithm at the centre of the model.

#### **3.1 SDSS overview**

Three main criteria were selected as the basis of the SDSS' evaluation procedure. The model sought locations that were safe while also maximising potential demand for medical services and centrality of the site. These were all categorised as very important in the literature review, giving them strong grounds for inclusion. Criteria relating to land use, land ownership and utility networks were omitted. While these remained valid, they did not emerge as high priorities during the literature review and are probably best addressed in person by teams when they view potential sites on the ground.

Figure 3.1 is a schematic of the SDSS, divided into three distinct phases. At its core was a NSGA-II genetic algorithm, chosen because of its ability to evaluate two criteria (in this case population and centrality). It was preceded by a setup phase and safety phase that prepared the candidate locations for evaluation and filtered out any at dangerous locations. Each of these phases involve their own spatial analysis processes, which will be described subsequently. Aside from definition of the candidate location grid, which was carried out in QGIS, all components were written in Python using Anaconda Jupyter's IPython kernel.



**Figure 3.1:** Main components of the SDSS divided into three distinct phases.

## **3.2 The setup phase**

The setup phase defined the candidate locations and calculated their scores for patient population and centrality to prepare them for evaluation later in the SDSS. This phase did not occur every time the SDSS was run and was considered as a data preparation stage that only needed to be rerun if any of the four inputted data layers changed. These layers and their evaluation will now be expanded upon.

### **3.2.1 Defining candidate locations**

As heuristic optimisation search methods, genetic algorithms require an initial population of candidates to determine a solution. One of the tensions here is selecting a population large enough to have sufficient diversity but not so big that the algorithm requires an inordinate amount of computing power and time to run (Diaz-Gomez and Hougen, 2007). Determining the optimal population size becomes more challenging when the complexity of the problem increases and in some cases a definitive best size may take considerable effort to quantify (Piszcz and Soule, 2006).

Before experimenting with population sizes, a method needed to be devised to create a pool of initial candidates. One option was to randomly generate coordinates in the study area, with these undergoing a brief screening to filter out any that are completely unsuitable, such as a point in the middle of a lake. This crude solution would work in theory but a big range of suitability could be expected across candidates at the outset, meaning more time would be required for the algorithm to return an acceptable solution.

Instead, a better approach was to leverage the 20 mobile hospital candidate locations provided by the WHO for the Mosul case study. Although there was no information on why they were originally selected, an analysis of common characteristics could help develop filters for screening the new pool of candidate locations, improving the overall suitability of these candidates at the outset. Taking this argument into account, QGIS was used to analyse the supplied candidate locations and then create a filtered candidate location grid. Note that some CHEs may not have the luxury of a pre-existing candidate list so if the SDSS is being adapted to a different scenario this step will need modification. A full explanation of the results of the analysis and the generation of the candidate grid is provided in the model development chapter.



### **3.3.2 Calculating patient population and TSPs**

After defining the candidate location pool, each point underwent assessment to determine the quantity of the nearby population and TSPs.

#### **Reason for inclusion**

For the mobile hospital to be effective, it needed to be at a location where it could serve the greatest number of patients. This was calculated for each candidate location ahead of the genetic algorithm phase which evaluated this data. Two general demand cases were identified. Firstly, the everyday healthcare needs of the population, which could include displaced people and the host population depending on user preference. Secondly, victims of traumatic injury who arrived at TSPs and depended upon referral to a more sophisticated healthcare facility for treatment.

#### **Assessment**

The first user case took the population recorded as living within an acceptable and realistic driving time of each candidate location. Literature on mobile and field hospitals emphasises emergency care provision rather than general healthcare meaning that there is minimal discussion on the appropriate size and travel times for the latter type of service area. In lieu of this, the SDSS took the population within a driving distance of two hours, but this was easy to modify. The score for this first user case was evaluated by the GA later in the SDSS.

For the second user case, the number of TSPs within an hour's driving distance of each candidate location were calculated, based on the 'golden hour' concept. A conservative measure of 30km/hr was used to calculate driving distance for both user cases following concerns about risky driving behaviour to meet the 'golden hour' raised by Rogers, Rittenhouse and Gross (2015). The score for this second user case was not evaluated by the GA, but was shown to the user as supplementary information during visualisation to help them make a decision on the final location.

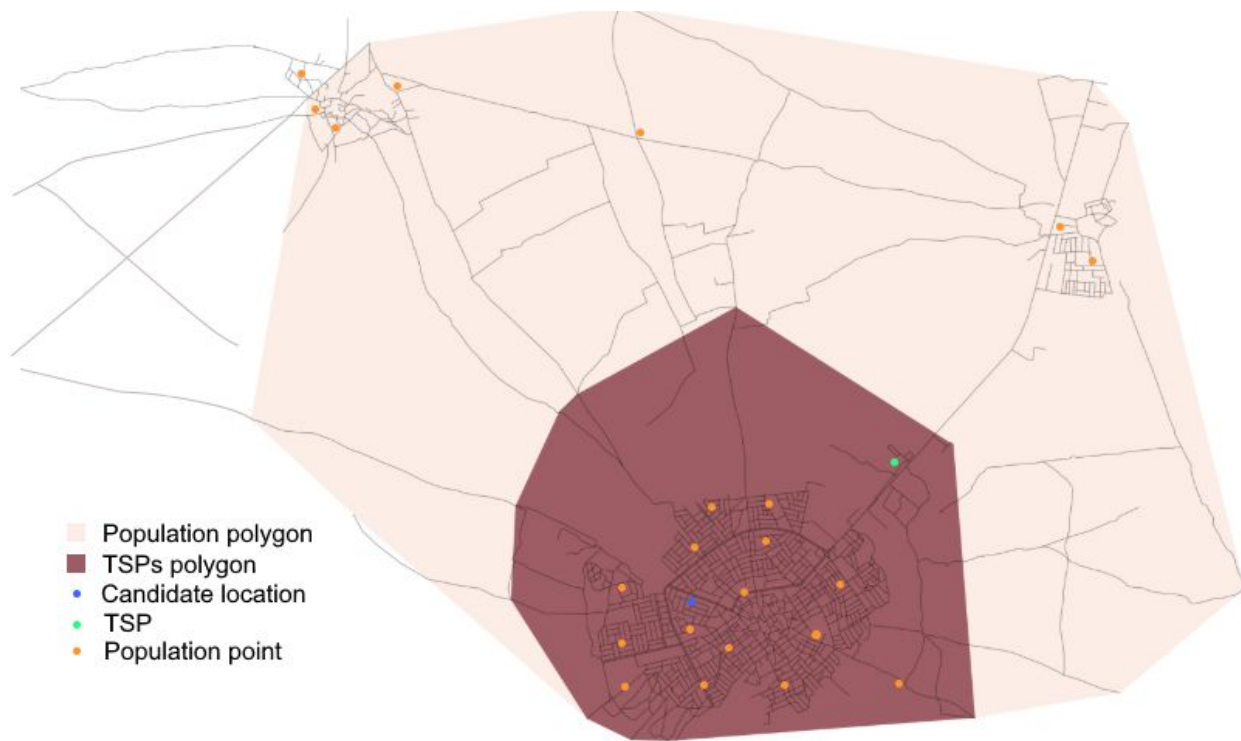
#### **Methods**

Calculating driving distance required network analysis. The method used here drew inspiration from Shukla, Wickramasuriya, Miller and Perez's (2015) procedure for estimating the population coverage of radiotherapy centres in New South Wales, Australia. The OSMnx tool for Python was used to download the road network for the study area from OpenStreetMap. An identical procedure for all candidate locations was then followed, beginning by finding the nearest node on the network to each candidate. Using functions in the OSMnx library, all nodes within a driving distance of this nearest node were returned and used as parameters for the construction of polygons representing the reachable area from the candidate location's nearest node. This

was carried out twice, resulting in a polygon representing one hour's driving distance (the 'golden hour') and a polygon corresponding with two hour's driving distance.

The spatial relationship between the first polygon and points representing TSPs was then assessed. Any TSPs inside the polygon were counted and the sum of these was added to the candidate location table. In the mockup shown in Figure 3.2, the candidate location would have one TSP since there is only one inside the TSPs polygon.

The procedure for the second polygon was similar, but had an additional step. Its relationship with a population layer consisting of points with attribute data on numbers of people was assessed. The number of people for each point inside the polygon were summed together. This total sum was then added to the candidate location table. In the mockup shown in Figure 3.2, the candidate's score would be the sum of all the population counts of the points inside the population polygon.



**Figure 3.2:** Mockup example of drive time polygons using the OSMnx tool to analyse the street network of the town of Karakush and its hinterland. Note that the points are dummy data and drive times have been restricted to only a few minutes for demonstration purposes.

### **3.3.3 Calculating centrality**

The second assessment component in the setup phase derived a score for each candidate location's centrality.

#### **Reason for inclusion**

Good road access is necessary for the deployment of the mobile hospital to the site and for the arrival and departure of patients and supplies. In a CHE sections of the road network are likely to be impassable due to damage, checkpoints and fighting. To mitigate a site being cut off from the rest of the road network, locations with multiple access points were favoured.

#### **Assessment**

The degree centrality of the nearest node to each candidate location was used as an indicator of accessibility. Degree centrality refers to the number of edges (in this case roads) that connect to the node. A node at a dead end has a degree centrality of one whereas a node at a crossroads has a degree centrality of four. The latter node can be reached from multiple directions while the former node is vulnerable to being cut off from the rest of the network. This method was admittedly crude, but its simplicity meant it could be quickly calculated. It was also easy to understand which was a potential advantage in a tool pitched at decision makers who may not have a geoinformation background.

#### **Method**

OSMnx was used to calculate the degree centrality for all nodes on the network. This returned the number of edges (road segments) connected to each node in the network, and then normalised it by dividing it by the maximum possible degree expressed as a percentage. For each candidate location, the nearest node's degree centrality was added to the candidate location table.

### **3.3 The safety phase**

The second phase of the SDSS was designed to remove any candidates located in restricted zones. It also normalised the population and centrality scores of the remaining candidates so they were ready to be assessed by the GA. The user was able to run this phase whenever they wanted to change the date that corresponded with different restricted zones or input a minimum distance from these zones.

#### **3.3.1 Removing any candidates in dangerous locations**

##### **Reason for inclusion**

CHEs are extremely dangerous environments so it was vital to position mobile hospitals in places where both staff and patients are relatively safe. The fast moving nature of contemporary conflicts, possible presence of unaccountable armed militia groups and long range weapons made it unlikely that any location could wholly guarantee safety, but distance from the main areas of fighting was a reasonable start.

##### **Assessment**

Due to its importance this criteria was realised before the candidates were evaluated by the genetic algorithm. Each time the SDSS was run, users could choose a date to determine the restricted areas and then the euclidean distance of candidates to the nearest restricted zone was calculated. Under the default setting any candidates inside or within one metre of a restricted zone were excluded. Users were able to expand this by inputting an acceptable minimum distance from the restricted zone if they had additional intelligence that suggested areas nearby may also be dangerous. Note that straight line distance was chosen over network distance because ordnance and some hostile military vehicles do not need to rely on established roads to arrive at the mobile hospital location.

##### **Methods**

For every candidate location the euclidean distance to each restricted polygon was calculated then the smallest distance from this list was returned. Any candidate locations where this distance was smaller than one metre (default) or a user specified distance were then dropped from the table.

An additional function included in early versions of the SDSS allowed users to add a buffer of  $n$  distance around each candidate and then drop any with military installations inside this buffer. Spatial segregation of humanitarian agencies and military facilities is encouraged to avoid conflation of the two by enemy forces (and a potentially volatile population). This process was not fully implemented due to a lack of data for military installations in the case study.

### **3.3.2 Normalisation of population and centrality scores**

The remaining candidates formed the pool from which the genetic algorithm randomly selected its starting population. A similar scale for the scores that evaluated by the genetic algorithm was advised by the DEAP development team because of the way the crowding distance measure works (De Rainville, 2018). This can be achieved through weighting each objective when setting up the algorithm or by standardising the existing scores themselves. The former required an additional stage of testing to choose the best weighting combination and would depend heavily on the scores for the set of candidate locations whereas the latter could be quickly implemented in Python for each new set of candidate locations, so in this instance was preferable. A simple min-max normalisation method was used to convert the scores for the two variables to numbers between 0 and 1.

## **3.4 The genetic algorithm phase**

The nucleus of the SDSS was a NSGA-II algorithm that came included in the DEAP library. This algorithm was selected because it was able to evaluate individuals against multiple objectives, ultimately returning a generation of individuals that made up a pareto front. Although the algorithm came pre-written in the DEAP library, decisions still needed to be made about the evaluation criteria, mutation probability, crossover probability, termination conditions and population size. A series of systematic experiments were used to assess options for some of these parameters, the methods and results of which are described in the model development chapter.

### **3.4.1 Visualisation of solution**

The solution arrived on by the genetic algorithm was visualised in a new window using Folium. Attribute data of the candidate locations forming the solution was also shown along with the restricted area to give users additional context. This was the final stage of the SDSS.

### 3.5 Data

The strength of the SDSS' results were determined directly by the quality of the inputted data. As explored by Barakat and Ellis (1996) there are logistical and ethical challenges to gathering accurate data in a dynamic CHE environment. For the Mosul case study data collection was carried out in the field by the WHO and subsequently analysed by authors based at the WHO and University of Twente including Kakakhan et al. (2018). As such, there was a centralised library of data to draw upon for almost all of the key variables in the SDSS. Table 3.1 shows the data that contributed to the SDSS alongside the requirement(s) they satisfied.

**Table 3.1:** Data layers feeding into the genetic algorithm.

Requirement	Data	Source
Candidates to be evaluated.	Candidate mobile hospital locations.	WHO/University of Twente
Evaluating safety.	Restricted areas (militarised zones, ISIS held zones)	WHO/University of Twente
Evaluating potential patient population.	Population distribution, location of IDPs	WHO/University of Twente
	Location of TSPs	WHO/University of Twente
Evaluating centrality and potential patient population.	Road network	OpenStreetMap.

### 3.6 Tools

There were multiple software products available that could be used to write the algorithm and create the SDSS, ranging from ESRI's well established Arc suite to open source offerings. Where possible this project favoured open source tools and sought to maximise interoperability. In CHEs where resources are scarce and analytical work may be carried out by volunteers, it is important that barriers to entry are as low as possible. Software licenses can have a prohibitive impact on volunteer capacity and engagement (Leidig and Teeuw, 2015) but fortunately there is a burgeoning range of open source GIS options with strong capabilities (Barik, Samaddar and Gupta, 2009). In addition, an open source approach meant that this project could be replicated and expanded without the need for significant financial investment.

Unpacking the methodology and design rationale of existing SDSSs revealed some suitable tools, with these listed in Table 3.2. The PostgreSQL database management system with the PostGIS extension is a popular spatial data repository because it allows for efficient storage and

SQL querying, as demonstrated by Modica et al. (2016) and Mekonnen and Gorsevski (2015). For sourcing and carrying out network analysis of OSM data, Boeing (2017) devised an elegant Python package that allows users to run complex operations with just a few lines of code. Python was also attractive for writing the genetic algorithm itself because of the availability of prebuilt frameworks, among which the DEAP library was well reviewed (Kim and Yoo, 2019) and had been used to solve various multiple criteria location allocation problems (Toutouh, Rossit and Nesmachnow, 2019; Zeng et al., 2019). Lastly, Folium powered the visualisation components of the SDSS while Tkinter was used to create a basic graphical user interface (GUI).

**Table 3.2:** Main tools and software for creating the SDSS.

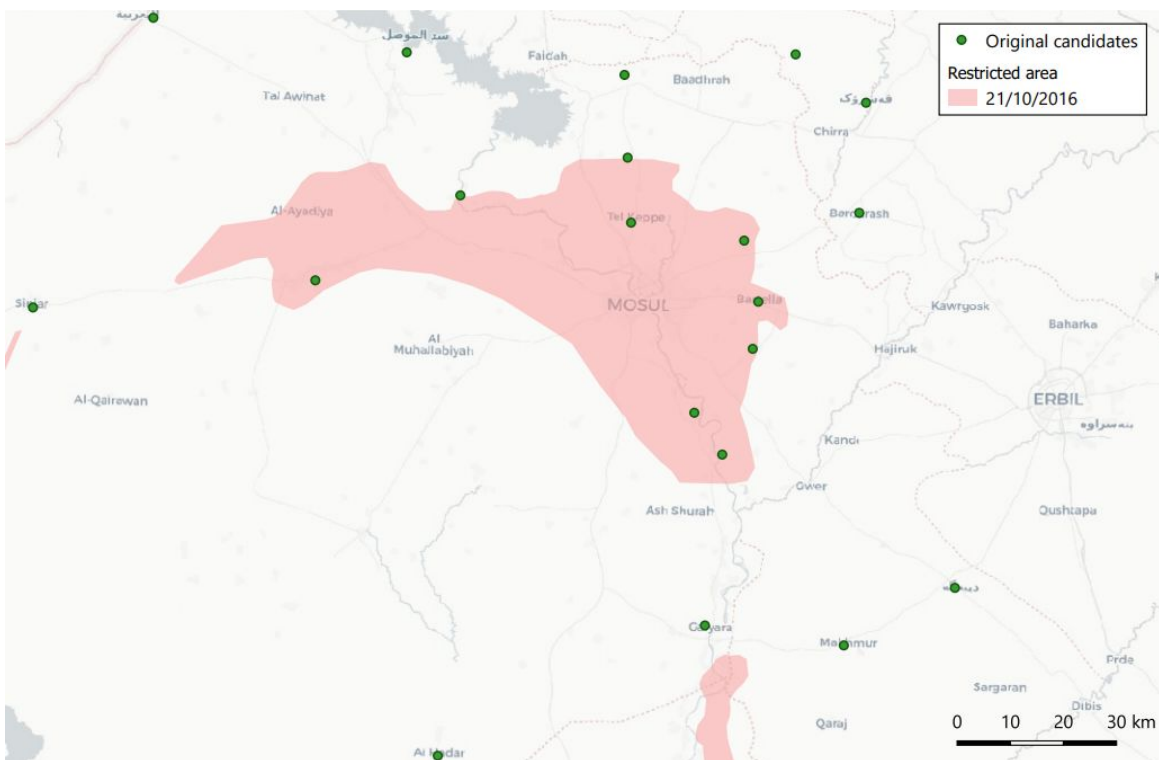
<b>Tool</b>	<b>Reasoning</b>
Python	Programming language for writing the genetic algorithm and evaluation criteria functions.
DEAP	Python framework for constructing genetic algorithms.
OSMnx	Python package for downloading OpenStreetMap road network and performing network analysis on it.
Geopandas	Python package for organising spatial data.
Tkinter	Python package for creating a graphical user interface (GUI).
QGIS	GIS software for analysing candidate locations and organising input data layers.
Postgres and PostGIS	SQL database for storing all data layers apart from road network.
Folium	Visualisation of solution delivered by genetic algorithm.

## 4. Data preparation and model development

Having described the SDSS in general terms in the previous chapter, it is now time to explain the steps that were taken to apply it to the Mosul case study and establish its key parameters. This chapter covers these items, which include the generation of the candidate location grid, establishment of an effective population size and termination generation for the genetic algorithm, and selection of the best individual solution from the pareto front.

### 4.1 Analysis of supplied candidate locations and generation of new locations

An analysis of the 20 candidate locations used during the initial WHO project was carried out to identify commonalities with the intention that this information could then be used to generate additional candidates. The analysis looked at how the candidates related to the road network, each other, Mosul city and restricted areas. In the supplied data restricted area polygons were tagged with one of four timesteps to represent repositioning of forces in the CHE: 21 October 2016, 3 November 2016, 24 November 2016 and 9 January 2016. Figures 4.1 to 4.4 display the candidate points in relation to the changing restricted areas as territory shifted between forces during the battle.



**Figure 4.1:** Candidate locations and areas marked as restricted on 21/10/2016.



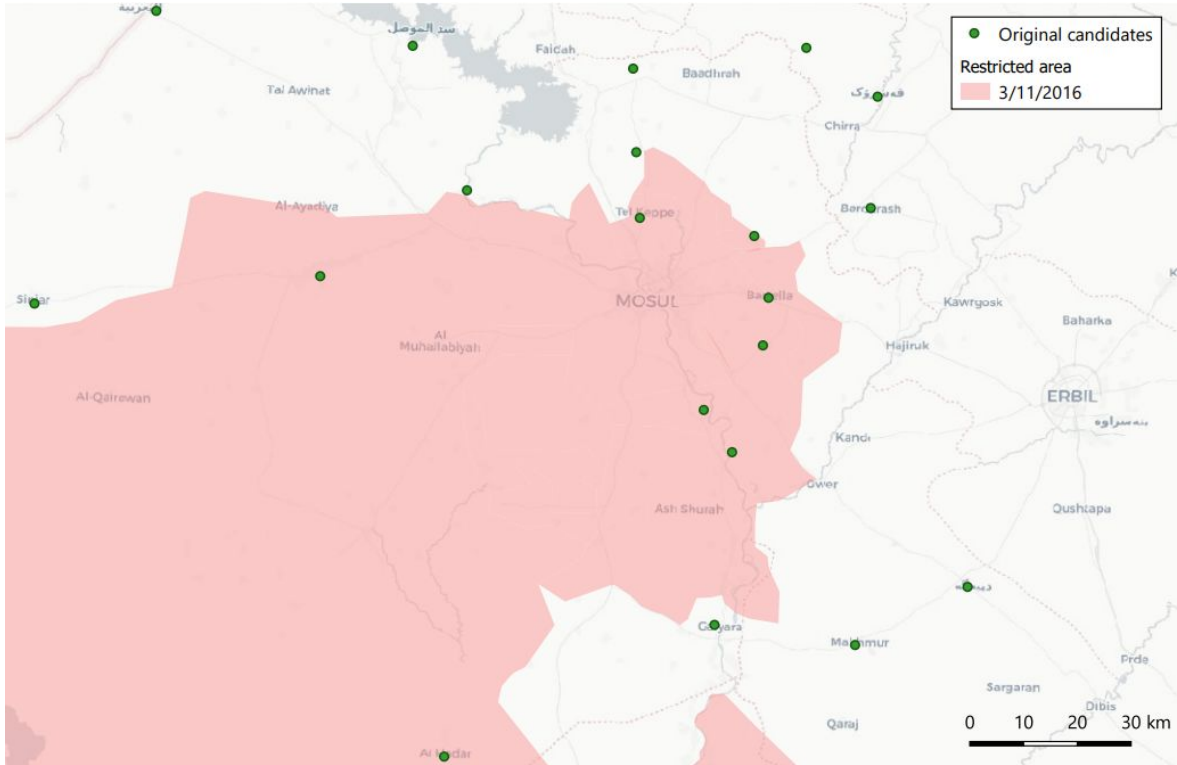


Figure 4.2: Candidate locations and areas marked as restricted on 3/11/2016.

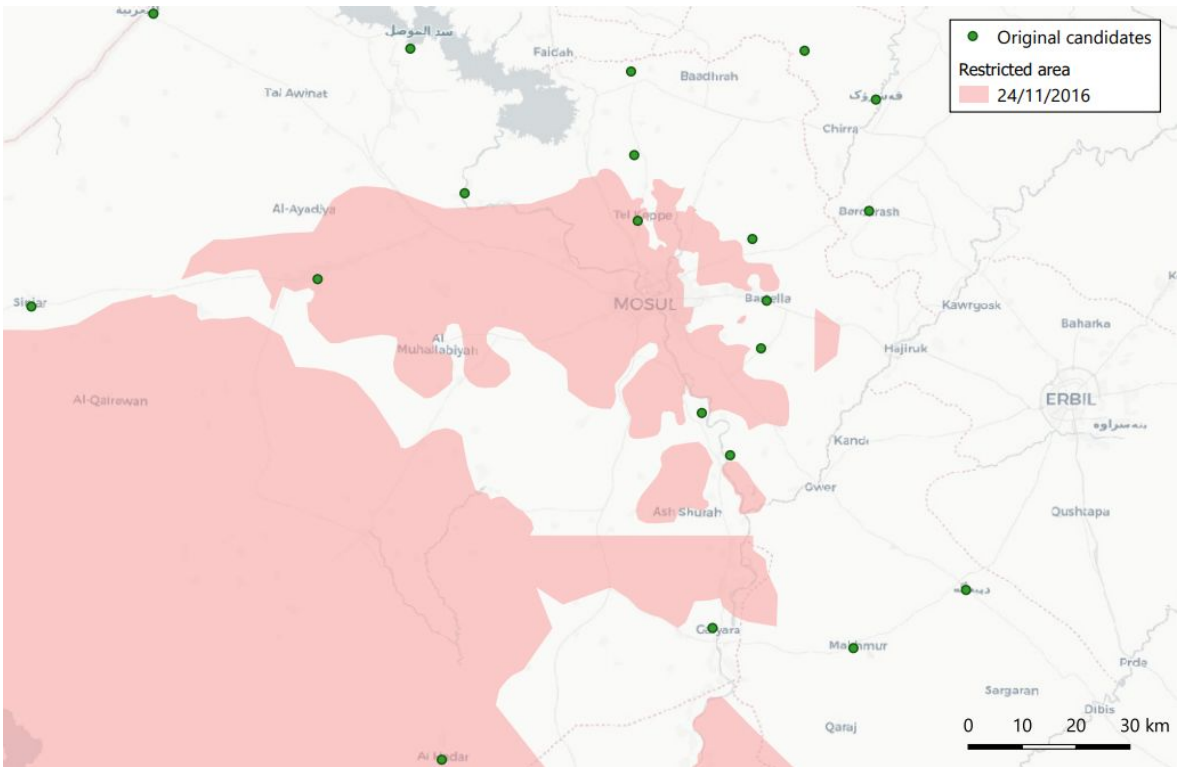
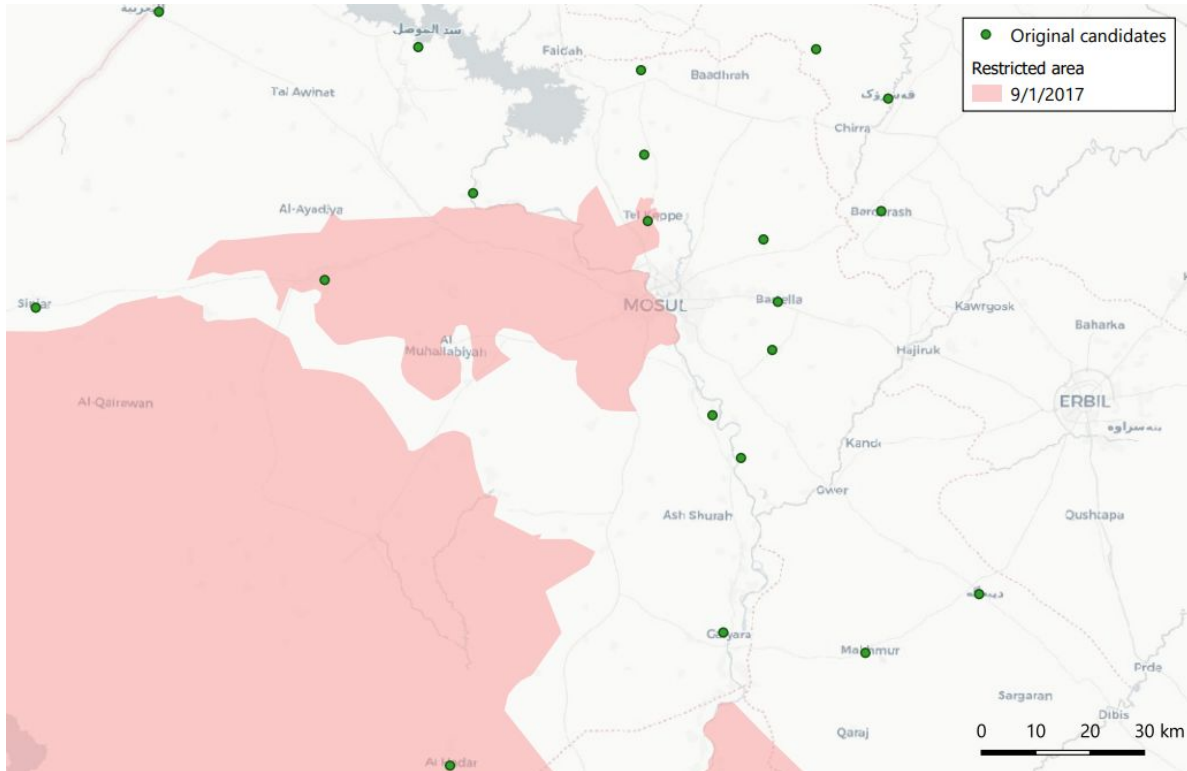


Figure 4.3: Candidate locations and areas marked as restricted on 24/11/2016.



**Figure 4.4:** Candidate locations and areas marked as restricted on 9/1/2016.

Table 4.1 displays the main findings from the analysis of the candidate locations. These were:

- Proximity to the nearest road ranged from immediately adjacent (<0.01km) to 1.6km, with candidates an average of 0.26km away.
- As seen in Figures 4.1 - 4.4, some candidates were situated in restricted zones with these numbers decreasing as ISIS retreated over the battle and more territory became available to coalition forces.
- No candidates were in Mosul itself, with the closest site being at Tilkef, located approximately 15.08km away from Old Bridge in the centre of the city. At the other extreme, candidate 11 at Sinjar was furthest away from the city, being 114.82km to the west of Old Bridge.
- Candidates were quite spatially dispersed. The closest two locations were 9.5km away from each other.

**Table 4.1:** Findings from analysis of candidate locations.

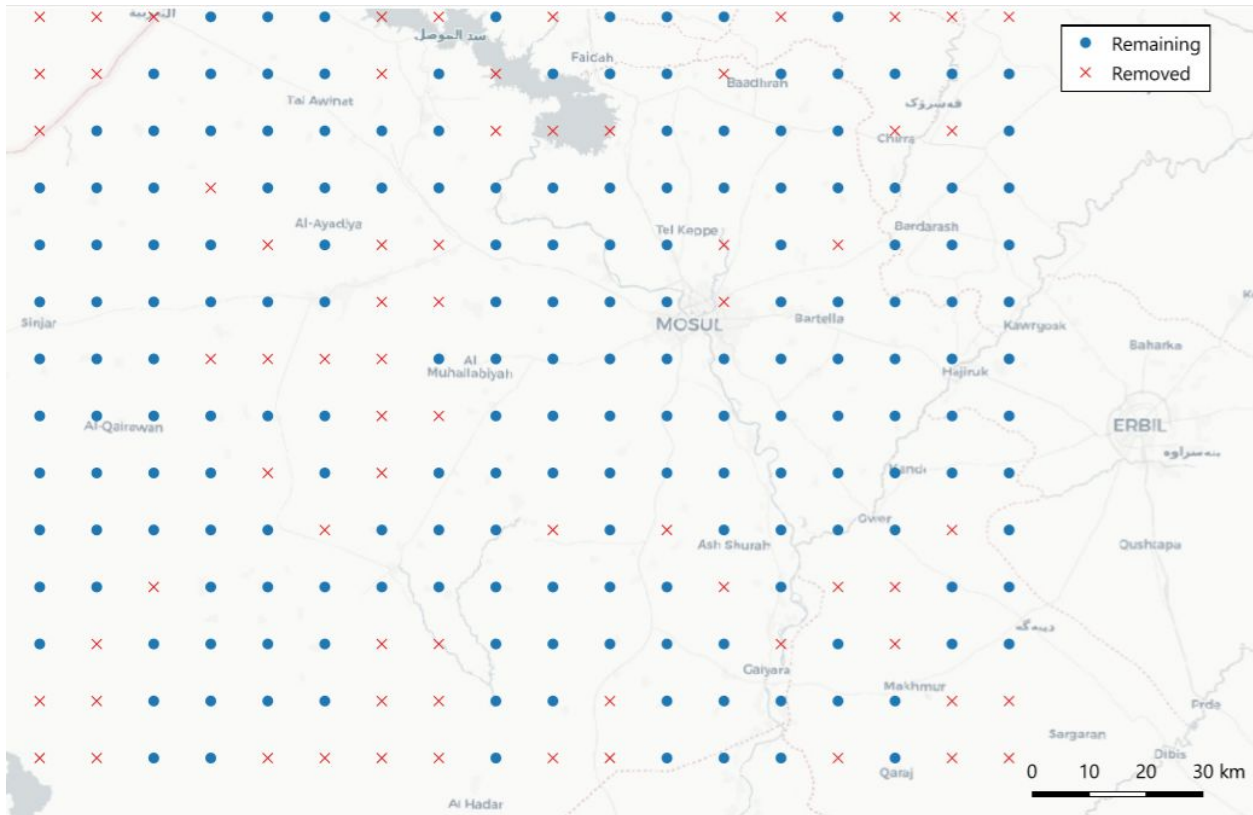
ID	Nearest town	Distance to road (km)	Distance to restricted zones (km)				Distance to Old Bridge, Mosul (km)
			21 Oct 2016	3 Nov 2016	24 Nov 2016	9 Jan 2017	
1	Bardarash	0.07	18.91	14.58	20.61	40.67	44.06
2	Al Hamdaniyah	0.01	Inside	Inside	1.89	17.16	22.48
3	Bartalah	0.09	Inside	Inside	1.71	18.88	21.73
4	Namrud	0.2	Inside	Inside	1.09	20.89	32.34
5	Hatra	0.01	48.75	Inside	Inside	Inside	93.64
6	Al Qayyarah	0.54	6.58	4.25	4.25	12.79	61.92
7	Ba'Shiqa	0.03	Inside	Inside	2.31	19.15	22.29
8	Hamam al 'Alii	0.12	Inside	Inside	1.41	13.82	22.96
9	Qasrok	0.31	30.18	30.48	36.25	46.93	56.30
10	Shekhan	0.33	30.09	31.31	33.72	40.68	54.72
11	Sinjar	0.18	4.73	4.37	4.37	4.37	114.82
12	Rabea' Area	0.09	41.07	34.71	43.52	43.52	106.53
13	Talafar	0.04	Inside	Inside	Inside	Inside	61.81
14	Zummar	1.6	21.44	27.86	26.25	29.68	64.65
15	Alqosh	0.01	15.76	15.08	18.58	22.73	42.83
16	Tal Usquf	0.46	0.23	1.47	4.83	7.83	27.29
17	Tilkef	0.39	Inside	Inside	Inside	Inside	15.08
18	Wana	0.57	0.35	0.94	2.64	2.64	39.79
19	Dibaga	0.06	41.45	34.46	34.99	37.51	79.64
20	Makhmur	0.29	18.49	14.81	14.81	23.68	74.83

#### 4.1.1 Generation of candidate location grid

Based on the findings of this analysis, QGIS was used to generate a grid of points at 10km intervals in a bounding box defined by the pre-supplied candidate locations. The grid method was chosen instead of random points within the bounding box to prevent spatial clustering of candidates. The grid points were then filtered to exclude:

- Those over 1.6km away from a road.
- Those inside the city of Mosul.
- Any in bodies of water.
- Any outside Iraq.

Figure 4.5 illustrates the full grid, distinguishing between removed and remaining candidates. Of the 252 points on the grid 69 were excluded while 183 remained.



**Figure 4.5:** Grid of new potential candidates, separated between points that were kept following initial filtering and those that were removed.

## **4.2 Calculating scores for evaluation**

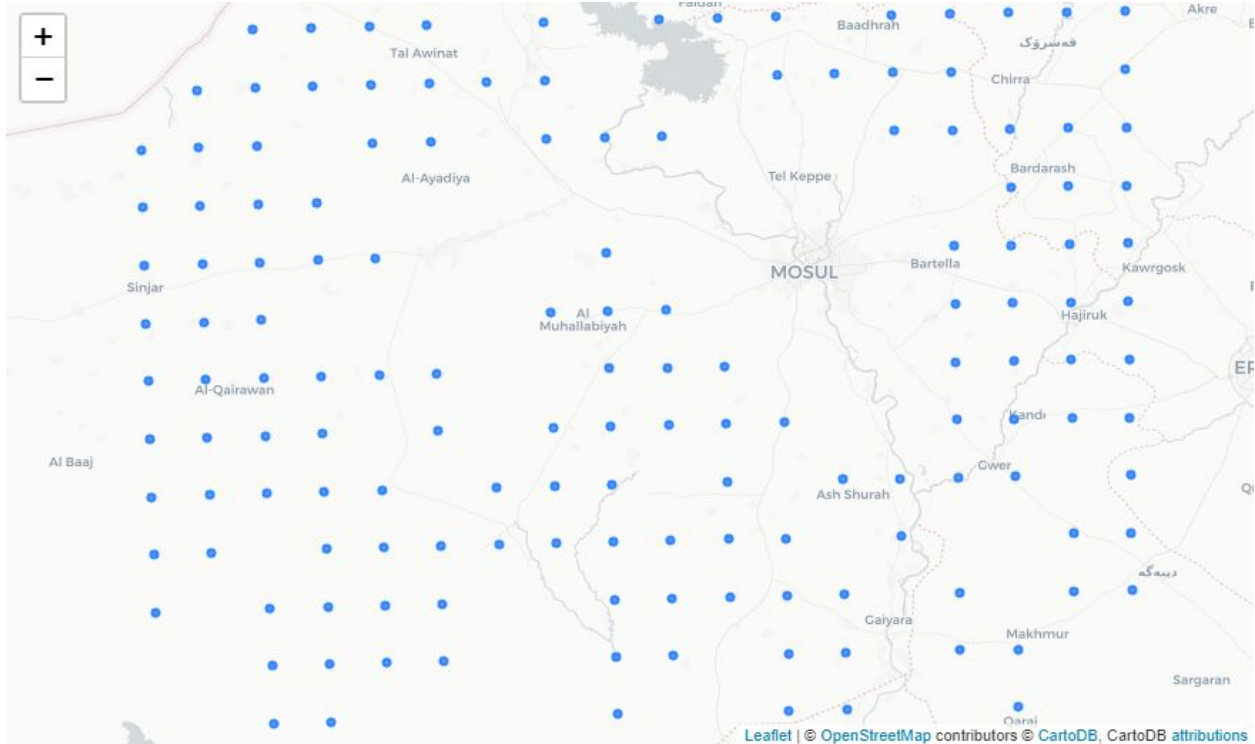
The grid was now ready to be worked on in a pure Python environment. Using the procedures described in the methodology, the setup phase was completed by calculating the patient population and centrality for each candidate location. To derive these scores population and TSP data layers were added from PostGIS while the network of drivable roads was downloaded from OpenStreetMap using the grid of candidate locations expanded by an additional 10km as a bounding box. This spatial extent could be modified quite easily by users, although it should be noted that downloading and projecting the network are two of the more system intensive functions in the workflow. During the model development IDPs outside camps were used as the population metric, but it was a straightforward procedure to change this for other measures of population (host population and IDPs inside camps being two other available sets of data).

## **4.3 Implementing the safety phase**

The SDSS was designed so the user could select the appropriate date to call restricted areas from PostGIS and then remove any candidates within a specified distance of these restricted areas. During model development the 21 October 2016 timestep was used to identify and drop any candidates inside the restricted polygons or within the default distance of one metre from the polygon.

### **4.3.1 Visualisation of candidate locations**

The 21 October restricted area safety filters removed a further 26 points from the already filtered grid of candidates. Figure 4.6 shows a visualisation in Folium of the remaining points. Note the spaces in the grid in the north-west corner (as this is on the Syrian side of the border), around Mosul itself and to the west of the city as this constitutes the main restricted area for October 2016. Other gaps in the grid correspond with two smaller restricted areas and locations where the road network was more spatially dispersed, so points were not within the 1.6km threshold of a drivable road.



**Figure 4.6:** Candidate locations visualised in Folium following 21 October safety filtering process.

**4.3.2 Normalisation and final candidate table**

Demand population scores across candidates ranged from 0 to 410,868 due to the large number of displaced persons in the Mosul theatre of conflict. The methodology for calculating centrality meant that scores ranged from 0.000024 (for nodes with only one connection) to 0.000097 (for nodes with four edges). These were both normalised to a 0 to 1 scale.

Figure 4.7 shows some candidates after the scoring functions and normalisation process were implemented. For evaluation by the genetic algorithm and subsequent visualisation, the main inputs will be the `idps_norm`, `cent_norm` and `tsp_count` columns.

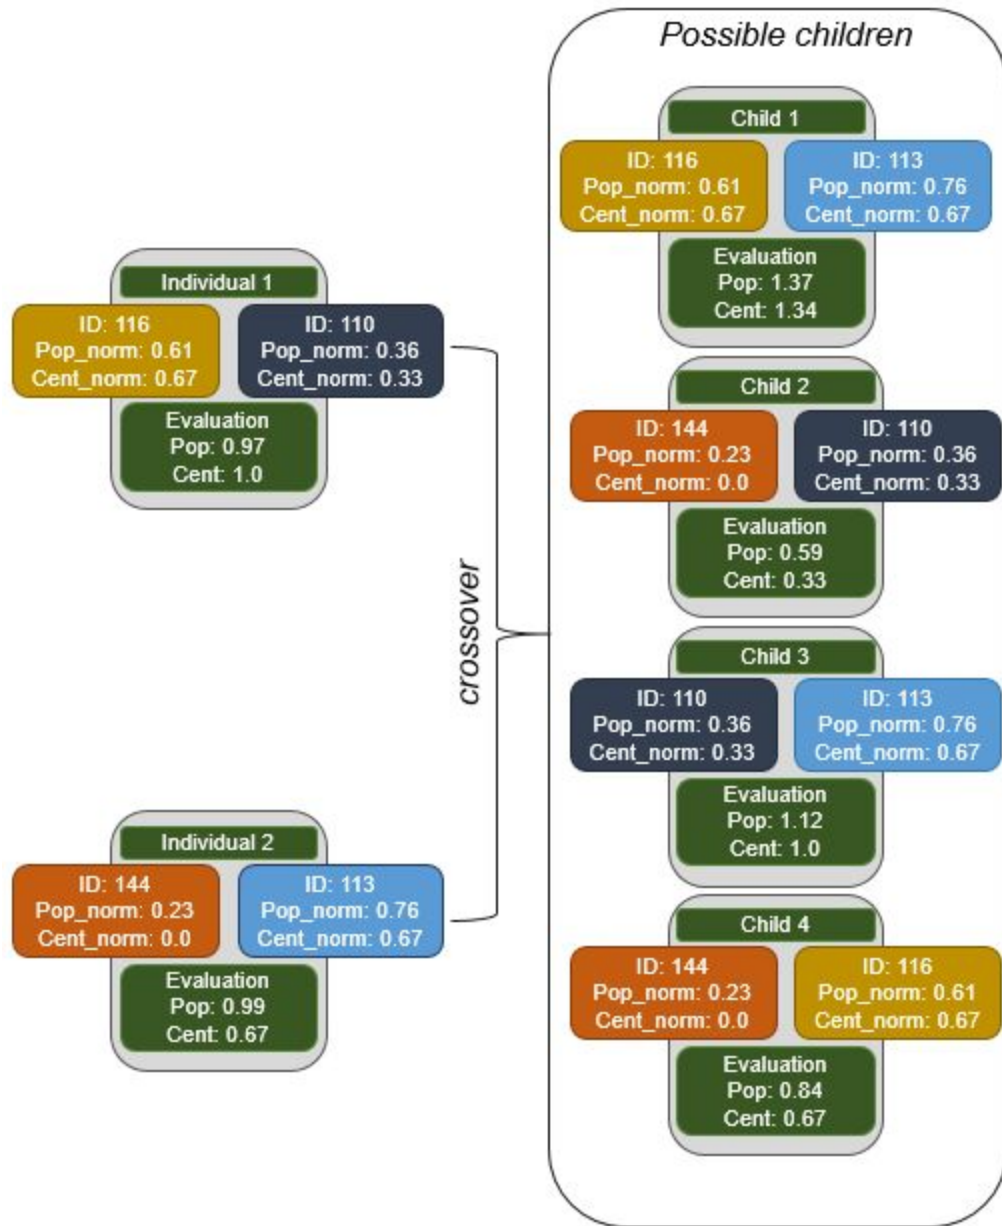
id	lon	lat	nearest_no	tsp_count	cent	cent_norm	pop_sum	pop_norm	geometry
1	218068.4864	4.047365e+06	4342138769	0	0.000073	0.666667	179555	0.047957	POINT (218068.486 4047365.418)
2	218068.4864	4.037365e+06	3022344920	0	0.000024	0.000000	161171	0.041391	POINT (218068.486 4037365.418)
3	218068.4864	4.027365e+06	3016828532	0	0.000073	0.666667	139119	0.033515	POINT (218068.486 4027365.418)
4	218068.4864	4.017365e+06	2993679190	0	0.000073	0.666667	95679	0.017999	POINT (218068.486 4017365.418)
5	218068.4864	4.007365e+06	3177340165	0	0.000073	0.666667	87314	0.015011	POINT (218068.486 4007365.418)

**Figure 4.7:** Five candidates with scores ready for evaluation by the genetic algorithm.

#### 4.4 The genetic algorithm phase

With the candidate locations now filtered for safety and scored for population and centrality, they were ready for evaluation by the genetic algorithm. For every run of the SDSS the algorithm created a population of individuals which consisted of pairs of candidate locations randomly selected from the pool defined in the safety phase. A pair of candidate locations was required for each individual so the algorithm could perform crossovers and produce children. Figure 4.8 illustrates this mechanic: if the two starting individuals were only made up of one candidate each, no crossover would be possible and the algorithm would not be able to produce new generations of stronger individuals. In this example, both Child 1 and Child 3 outperform their parents. This section of the model development explains the tests used to establish a suitable starting population size, termination generation and selection of the best individual.





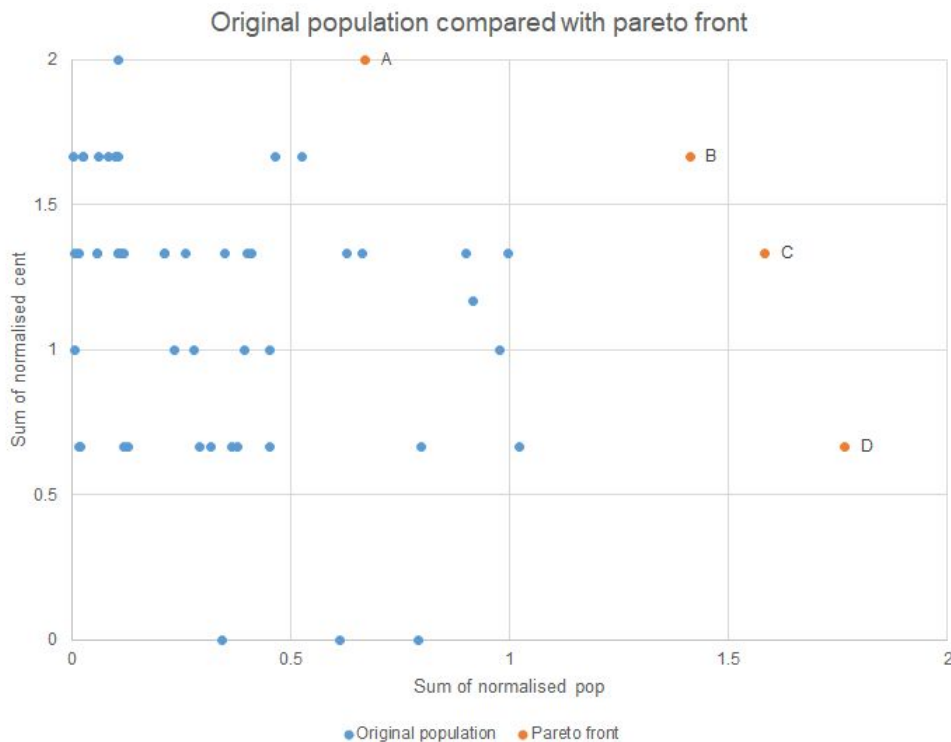
**Figure 4.8:** Crossover of two individuals and all their possible children.



#### 4.4.1 Genetic algorithm output

To gain insight into the genetic algorithm's output before testing, it was first run for 50 generations using a starting population of 50 individuals. The result was a generation of individuals forming a pareto front. Figure 4.9 visualises this by showing the scores for the original population of 50 individuals alongside the scores for those forming the pareto front after the algorithm had run for 50 generations. The pareto front was comprised of pairs of locations that had the highest cumulative normalised cent and cumulative normalised population scores, with IDPs outside camps used for the population layer.

In this instance the front contained an individual that maximised the centrality variable at the cost of the population variable (A), an individual that does the opposite (D), and two individuals that were more balanced between the two variables (B and C). All can technically be considered as ideal solutions so additional decision making was required to select the preferred individual from this front. While there was an evident improvement between the original population and pareto front, the graph also helps illustrate the importance of population size. By limiting the original population to 50, candidates that could have formed even better performing individuals than those in the pareto front may have been omitted. Both of these concerns - population size and selecting the best individual from the pareto front - will now be addressed.



**Figure 4.9:** Pareto front after running the algorithm for 50 generations with a starting population of 50 individuals.

#### **4.4.2 Crossover and mutation probabilities**

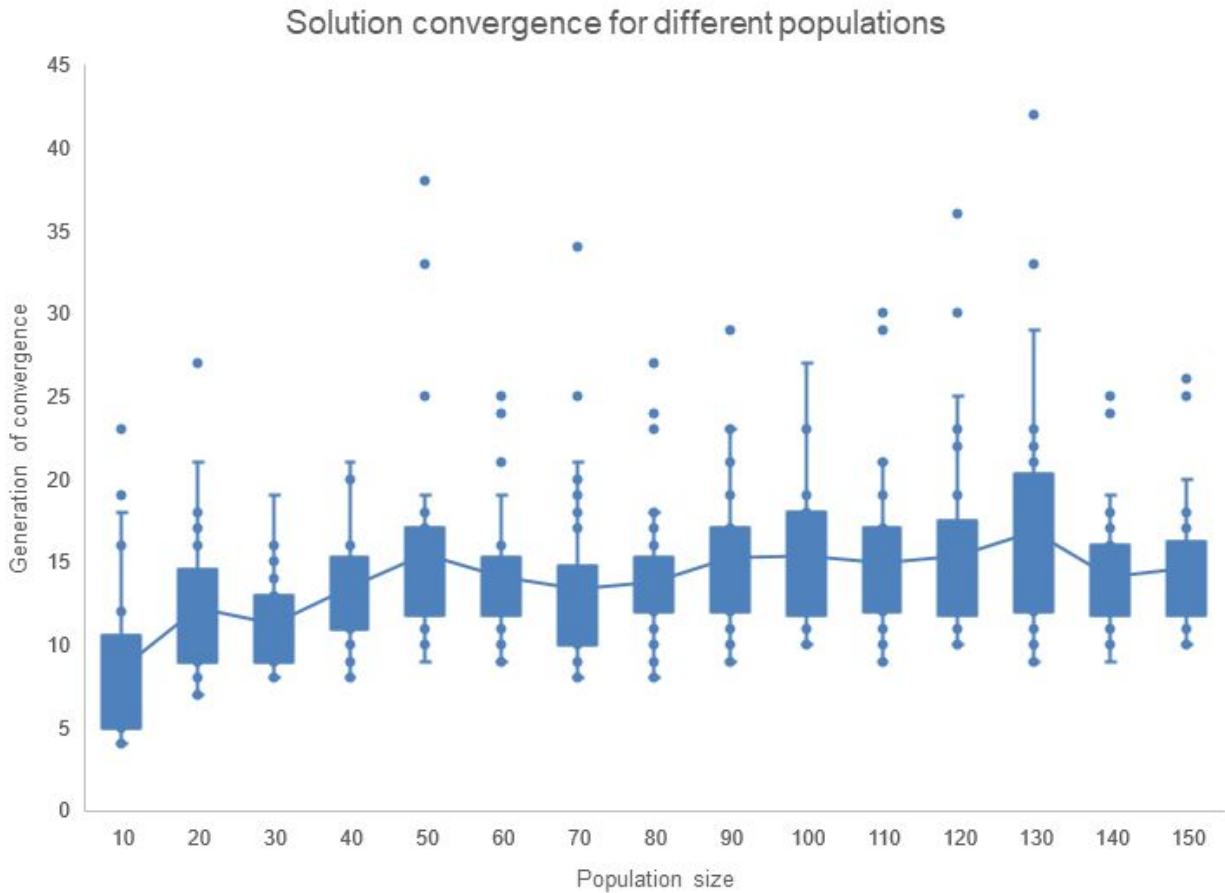
The mutation and crossover probabilities can also significantly influence the genetic algorithm's behaviour. A crossover probability of 0% means that no new offspring are created and the new generation is comprised of identical copies of chromosomes from the previous generation. A probability of 100% results in all offspring being created by crossover (Attia and Horacek, 2001). The crossover probability for the GA was set at 80%, based on the work of Deb et al. (2000). Experimentation with this parameter could be a direction for future work investigating suggestions such as Vasconcelos, Ramirez, Takahashi and Saldanha's (2001) adaptable approach or Attia and Horáček's (2001) changing probability rate.

An additional measure inside the evaluation function meant that any individuals made up of the same candidate locations were severely penalised so would not meet the threshold for staying in the generation. This step was included to prevent the algorithm from returning individuals made up of the same two candidate locations as a solution.

Mutation is a mechanism that can be written into a genetic algorithm to enhance and maintain diversity within the population, helping to uphold the algorithm's capacity for exploration of the search space and mitigating premature convergence. When mutation is included there is typically a probability describing how likely it is to occur for an offspring. A very high probability of mutation will prevent the algorithm from converging on a suitable solution, so the mutation probability for the genetic algorithm used in this SDSS was set at 20%. As with crossover probability there is opportunity for further experimentation.

#### **4.4.3 Population size**

Previous work on genetic algorithms demonstrated that larger population size offers the potential for greater genetic diversity and reaching a better solution. However, it also means that the algorithm requires more time to evaluate each generation and perform crossovers. To gauge a suitable size for the SDSS, the algorithm was run 30 times for populations between 10 and 150, with the population increased in increments of 10. As each individual in the population was made up of a pair of candidates randomly selected from the 157 candidates, a population of 150 individuals included a large sample, possibly even all, available candidates. The algorithm terminated after 50 generations for all runs.. There were two main questions that this test sought to answer: whether there was a significant change in the amount of generations it took to converge on the solution as population increased, and whether there was any variation in the quality of these solutions.

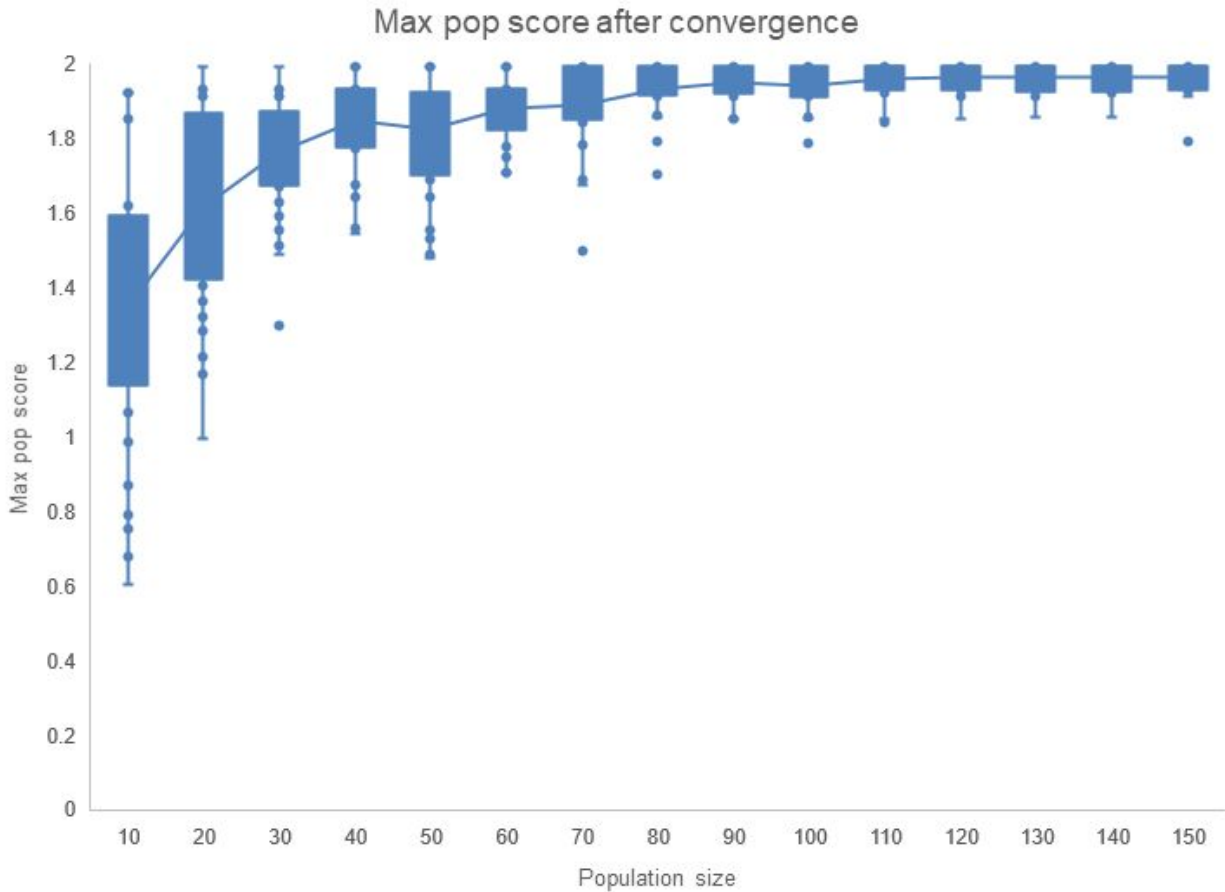


**Figure 4.10:** Identifying convergence for different population sizes (30 runs for each population size).

Figure 4.10 and its underlying data can help answer the first question. Regardless of population size, all runs of the algorithm converged in fewer than 50 generations, with the longest run taking 42 generations (for a run when population was set at 130). However, looking at the underlying data, convergence actually occurred after 12 generations - the makeup of each generation was stable until generation 41 when at least one individual suddenly changed. It was reasonable to assume that this change was caused by the mutation mechanism. This mechanic was significant for other runs with a relatively high generation of convergence. Among the 15 runs that took 27 or more generations to converge, 10 experienced a similar period of stability for multiple generations before a mutation happened. Although these instances should not be discounted, they are not related to the population size and can be considered as outliers.

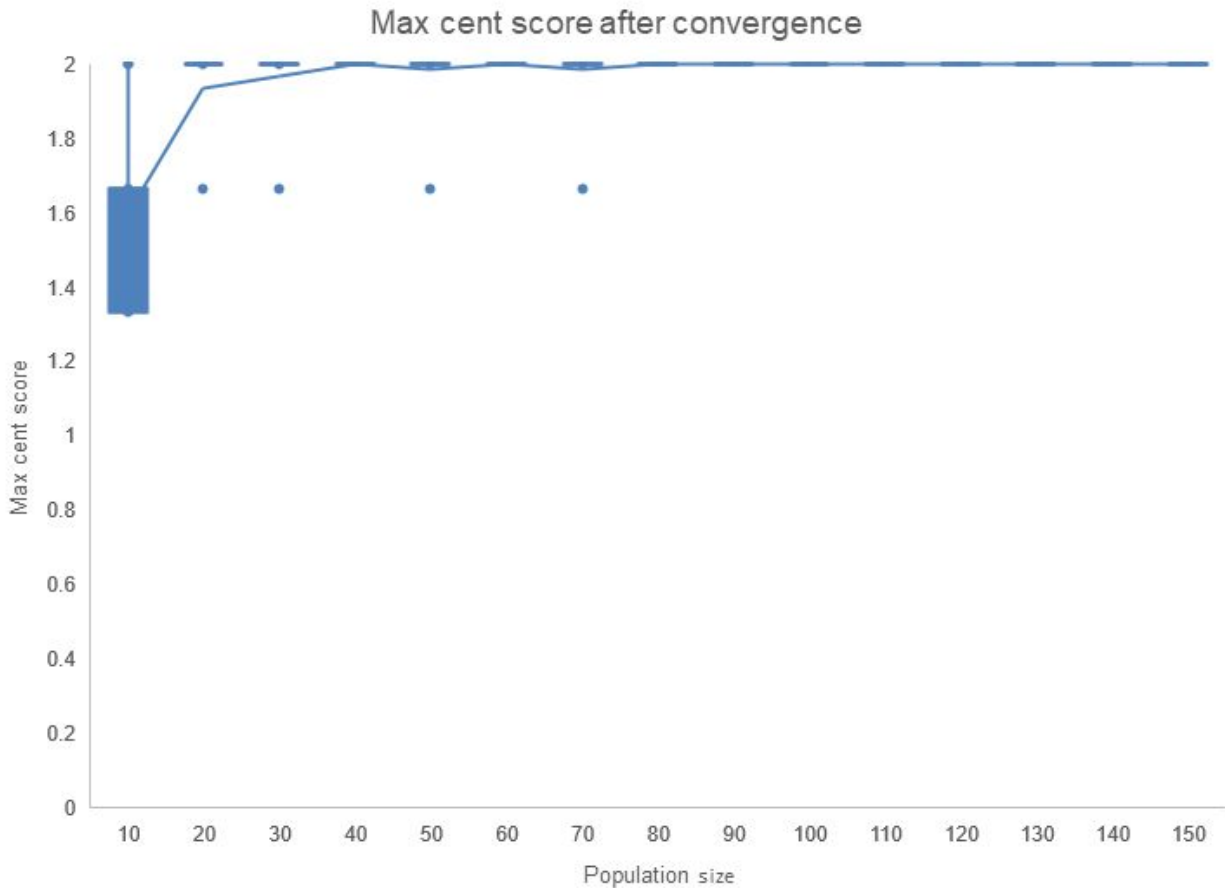
Most runs converged on a solution between the tenth and twentieth generations regardless of population size. The mean convergence generation across all runs was 14. Between population sizes 10 and 40 there was a small increase in the number of generations needed to converge, with the mean convergence generation for population sizes of 10 being generation 9. This behaviour was due to the limited crossover combinations between individuals. However, for this

data to really be useful the convergence generation needed to be considered alongside the quality of the solutions delivered. This was achieved by looking at the maximum scores produced for each variable by the algorithm over the various runs.



**Figure 4.11:** Maximum pop score among individuals for runs with different populations.

Figure 4.11 shows the way the maximum pop score changed when population was increased for each run. For these tests IDPs outside camps were the input layer used to derive the pop score. There was a clear relationship here: higher populations returned individuals with better performing pop scores more consistently than lower populations. This was expected because they had better genetic diversity and were more likely to contain candidates with the highest scores in their starting populations. While there was continued improvement as population increased, the shape of the graph suggests that a population of 80 or more will almost always deliver a maximum pop score of greater than 1.6.



**Figure 4.12:** Maximum centrality score among individuals for runs with different populations.

The results for centrality shown in Figure 4.12 corroborated the notion that bigger populations are more likely to reach higher scores, with 40 being the generation size where this stabilises. It should be noted that there was a greater skew towards maximum scores among the distribution of centrality scores among candidates than for the pop scores. This explains why even populations of only 10 were able to return individuals with the maximum possible centrality score and why the mean score line flattens out at a lower population than the one for pop scores.

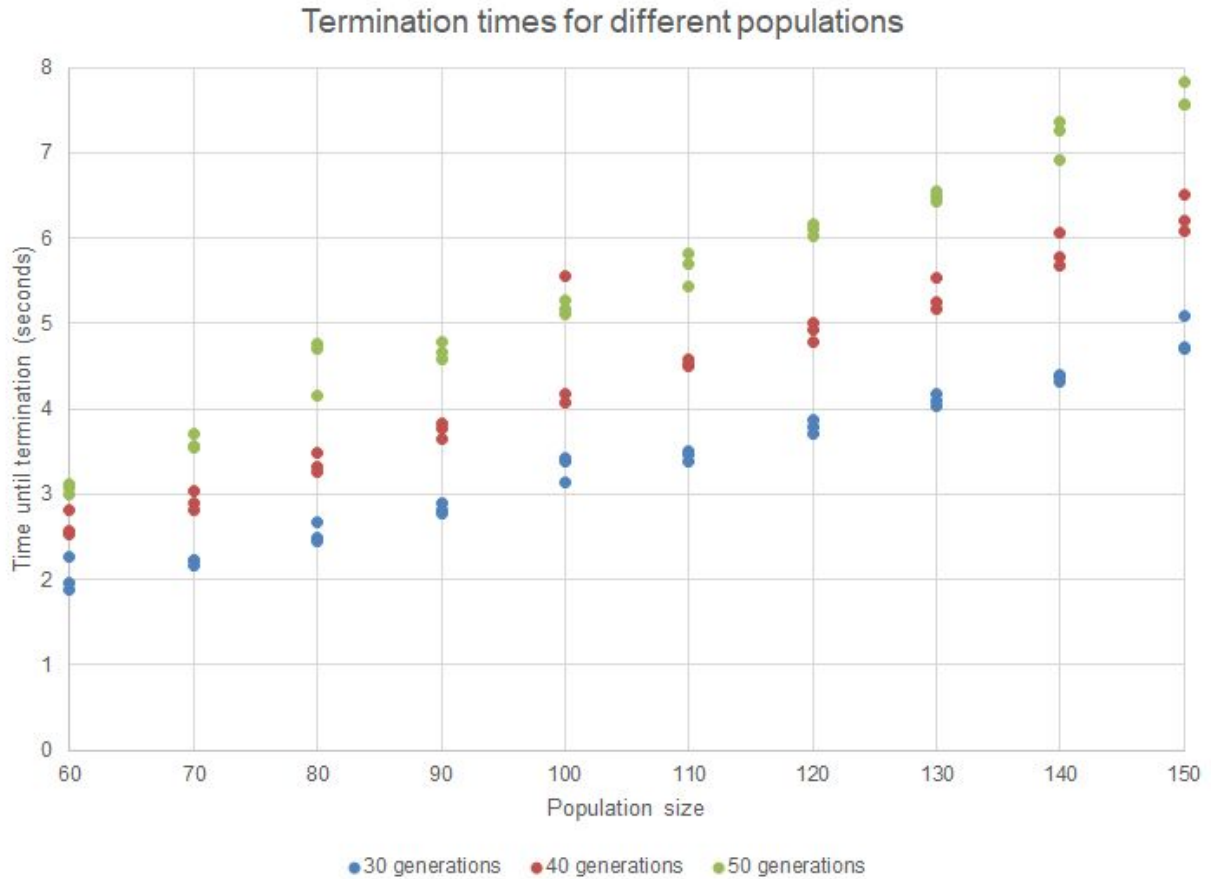
Taking the three graphs together, a strong case can be made for 80 being the minimum population size for delivery of a consistently strong solution. The algorithm is not expected to return the absolute best possible solution every time, but the maximum scores for the two variables are regularly in the upper quartile for solutions returned from populations of 80 or more. It is important to realise that the maximum pop score and maximum cent score belong to individuals at either end of the pareto front. A maximum pop score of 2 and a maximum centrality score of 2 does not in this case mean that there was an individual with the highest score for both variables. No two candidates in the dataset have perfect scores for both variables

so it was impossible for such an individual to emerge even through crossover and mutation. Experiments with datasets from other case studies may yield such an individual.

#### **4.4.4 Termination conditions**

Genetic algorithms are heuristics and as such are not designed to reach a definitive answer. They require an artificial condition to be satisfied in order to stop, meaning there is a design choice to be made about what this termination condition ought to be. This can be a maximum number of generations or when an individual emerges that satisfies pre-defined criteria. The former is appropriate when trying to find an individual with characteristics that are not already known (i.e. find the highest scoring individual, not find the individual with a score of 2). When collecting the data that was discussed in the previous section, the algorithm was terminated after 50 generations. For all population sizes a solution was reached well within this amount of iterations. The data indicated it would be possible to drop to 30 generations and still achieve strong results.

Figure 4.13 shows the time taken to reach termination for various population sizes and maximum generation counts. The results here add weight to the argument for setting the default termination generation to 30: among population sizes between 60 and 150, the algorithm was demonstrably faster at completing a run than when the termination generation was 40 or 50. However, caution should be exercised here. If the algorithm was repurposed for another scenario with a larger pool of candidate locations that in turn meant a larger population feeding into the algorithm, 30 generations might not be sufficient for convergence on a strong solution. In this case, of the four timesteps available to filter the candidate pool, the 21 October option used in these tests removed the fewest candidates, so there was no danger of the algorithm choosing a population from a larger pool if another timestep was selected.



**Figure 4.13:** Measuring time for the algorithm to reach completion under different population sizes and termination generations.

#### 4.4.5 Selecting the best individual from the pareto front

Having established a minimum population size that reliably delivered a strong solution, the next challenge was devising a method to select the individual from the solution that best satisfied the aims of the SDSS. Returning to the pareto front presented in Figure 4.9, it was easy to identify the individual with the best pop score and the individual with the best centrality score, but by maximising one variable the other suffered. To maximise both together, two individuals were contenders (B and C), forming the 'knee' of the pareto front. As Rachmawati and Srinivasan (2009) note, this convex bulge contains solutions that are typically optimal when it comes to the tradeoff between the two variables.

A systematic approach was required to meet the SDSS' objective of maximising both the pop and centrality scores. Because both variables operated on the same scale, a straightforward option was picking the individual where the difference between the two scores was closest to

zero. Table 4.2 demonstrates this, again using the front from Figure 4.4. With a difference of 0.248 individual C was closest to zero, but only marginally more so than B, which had a difference of -0.254.

**Table 4.2:** Differentiating between individuals in a pareto front.

Individual	Sum normalised pop	Sum normalised cent	Difference	TSP count
A	0.669047	2	-1.33095	2
B	1.412614	1.666667	-0.25405	5
C	1.582348	1.333333	0.249014	3
D	1.764742	0.666667	1.098075	3

#### 4.4.6 Final genetic algorithm parameters

Having completed the testing phase for the genetic algorithm, the final default parameters could now be defined. These were:

- Starting population of 80 individuals.
- Termination after the GA had evaluated 30 generations.
- Mutation probability kept at 20% and crossover probability kept at 80%.
- Return the individual from the pareto front which had the lowest difference between its pop and cent scores.

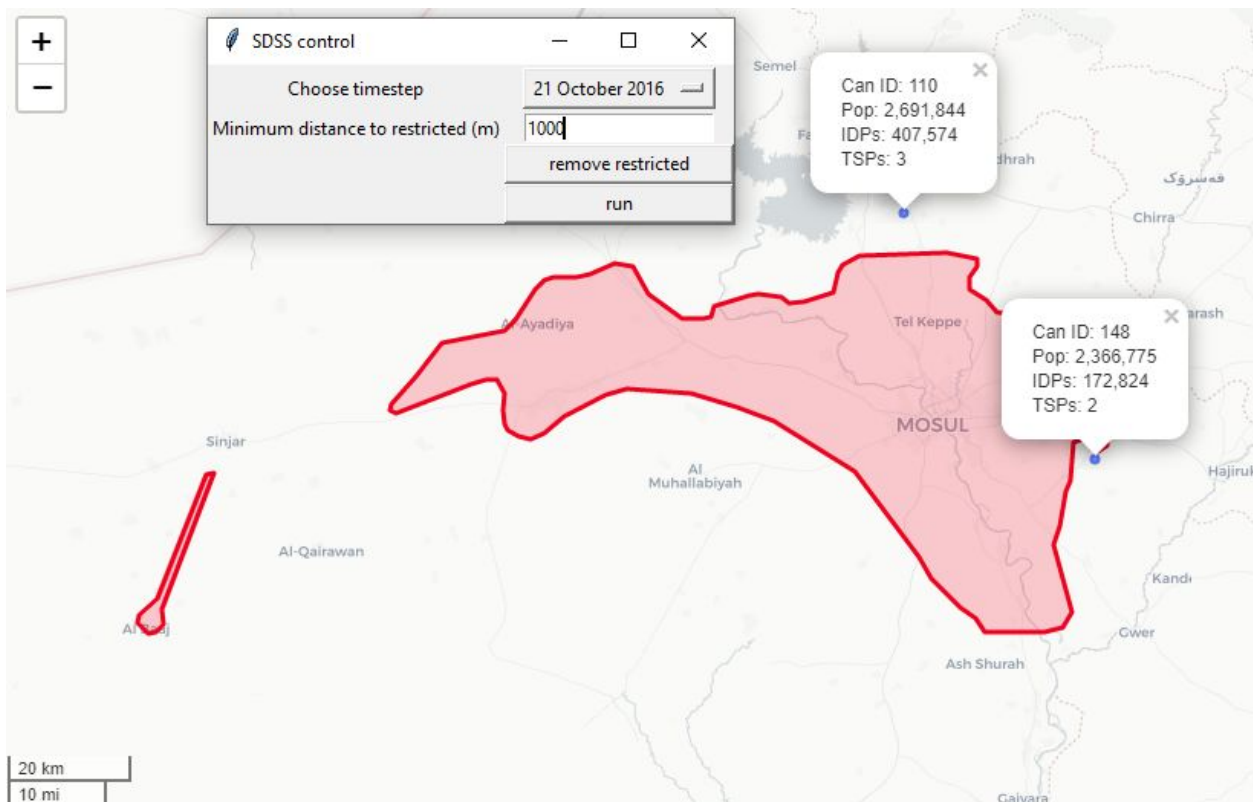
#### 4.5 User input and visualisation

The final development step was to package the safety and genetic algorithm phases together so a layperson could operate the SDSS. This was achieved by building a control dialogue window in Python's Tkinter package and through visualisation of the solution using Folium. Tkinter was chosen because it came included with Python 3 and is relatively easy to use. Ease of use was also a factor in selecting Folium over other mapping options, along with its wide range of built in basemap tilesets. Figure 4.14 shows the two elements together. By using the inputs on the control window, users were able to do the following:

- Choose between the four timesteps to toggle the different restricted areas.
- Input a minimum distance in metres for candidate locations to the closest restricted area. If this input box was left empty the default of one metre was used.
- Remove any candidates from the grid which fall inside the restricted areas and within the minimum distance set in the previous step.



- Run the algorithm. Each run opened a new browser window displaying the Folium map which centred on the study area and rendered an OpenStreetMap basemap in CartoDB's subtle Positron style.
- View the pair of candidate locations that made up the best performing individual returned by the algorithm. Popups show attribute information including the candidate ID, number of IDPs, total population and TSPs within the vicinity of the candidate.
- Either re-run the algorithm using the same input settings by pressing run again, or choosing a new timestep and minimum distance and removing restricted candidates accordingly before hitting run.




**Figure 4.14:** SDSS control window and model output, including restricted areas and two candidates making up the best individual returned after the run.

## 5. Results


### 5.1 Quality and temporal variation

This first item in the results section investigates the quality of the solutions returned by the algorithm and how these change according to shifting restricted areas and evaluation criteria. The intention was to determine the extent the results returned by the algorithm changed as the CHE unfolded, answering research question four. To generate results for interrogation, the algorithm was run 300 times for ten different combinations of evaluation criteria and restricted area. The evaluation criteria either sought to maximise centrality and the total population within a two hour drive, or maximise centrality and the number of IDPs outside camps within a two hour drive. The restricted areas varied for each of the four timesteps demarcating changes to the restricted areas, prefaced by a benchmarking exercise where no restricted areas were taken into account. Drawing on the conclusions of the model development phase, a starting population of 80 random individuals (pairs of candidates) was used for each of these runs, with the algorithm set to return a solution after 30 generations. The mutation probability was kept at 20% and the crossover probability kept at 80%. The results of these tests will now be presented in turn.

#### 5.1.1 Benchmarking exercise (no restricted areas)



**Figure 5.1:** Most frequently returned individuals after 300 runs evaluating centrality and total population (blue) and 300 runs evaluating centrality and IDPs outside camps (orange).



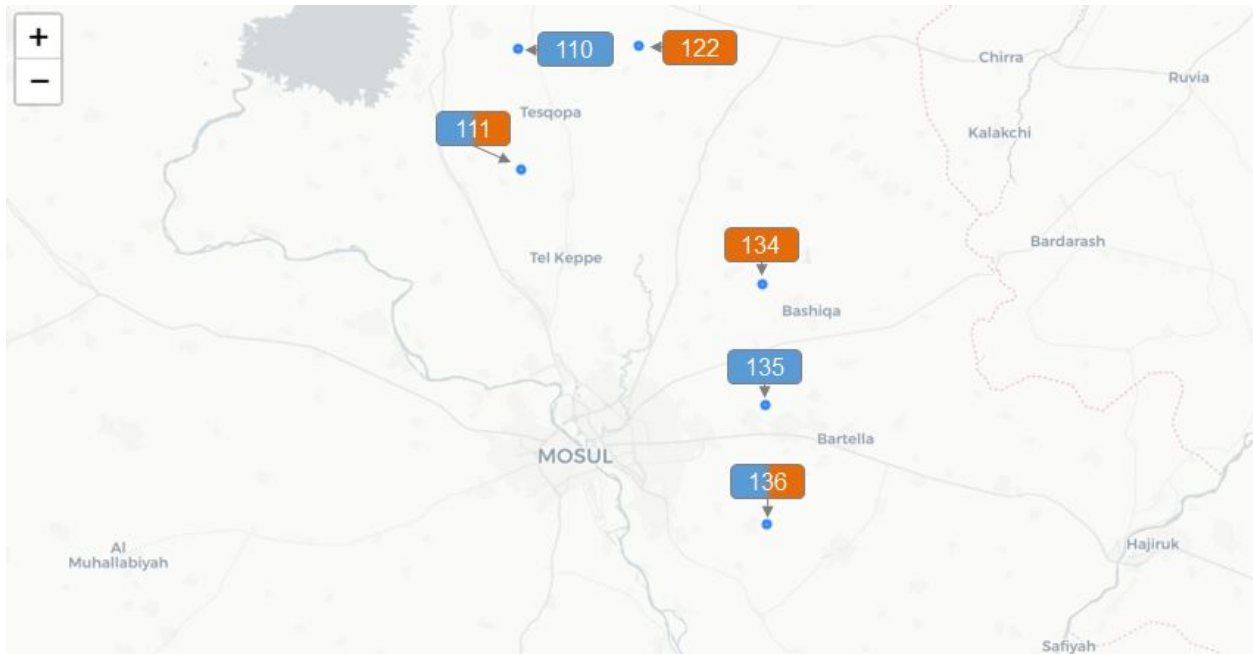
**Figure 5.2:** Normalised scores of the most frequently returned candidates after 300 runs evaluating centrality and total population (left) and 300 runs evaluating centrality and IDPs outside camps (right).

Figure 5.1 shows that the most commonly returned individual for the total population evaluation setting was made up of candidates 136 and 111, a pairing selected by 27% of the algorithm's 300 runs. Figure 5.2 shows why this was the case: among candidates with the maximum centrality score, candidate 136 had the highest population score, while candidate 111 had the highest population score within the pool of candidates scoring 0.66 for centrality. Within the overall candidate list the location which performed best overall for population (candidate 123) was situated at the end of a road so was given the lowest possible centrality score, explaining why it only featured in nine of the returned solutions.

111 also featured in the pairing most frequently returned for the IDPs evaluation, although for that test it was matched with 134 and was delivered in 31% of the 300 runs. Comparing the two graphs in Figure 5.2 shows that there is greater range variation among the normalised IDPs scores than the normalised population scores for candidates with maximum centrality. No candidate with maximum centrality achieved a score of over 0.71 for normalised IDPs while two candidates (136 and 134) scored over 0.93 for normalised population and had the maximum centrality score. This can be attributed to IDPs outside camps having a lower density and more fragmented distribution than the total population. Under both evaluation criterias, pairings of one candidate with the maximum centrality with another candidate of lower centrality but high normalised population or IDPs were a typical solution.

The visualisation in Figure 5.3 demonstrates that under both evaluation criterias the most popular candidates were all within relatively close proximity of Mosul, corresponding with the

high host population and IDPs in the periphery of the city. The impact of introducing restricted areas will now become apparent in the following subsections.




**Figure 5.3:** Visualisation of most frequently returned candidates for runs evaluating total population (blue labels) and IDPs outside camps only (orange labels). Half blue, half orange labels correspond with locations returned using both settings.


**Table 5.1:** Details of the most frequently returned candidates for the benchmarking exercise.

Candidate ID	Total population	IDPs outside camps	Degree centrality	TSPs
136	2,659,086	238,782	4	8
111	2,805,417	413,856	3	4
135	2,703,547	267,882	3	8
110	2,691,844	407,574	3	3
122	2,614,827	410,868	3	2
134	2,634,518	295,614	1	7

5.1.2 21 October 2016



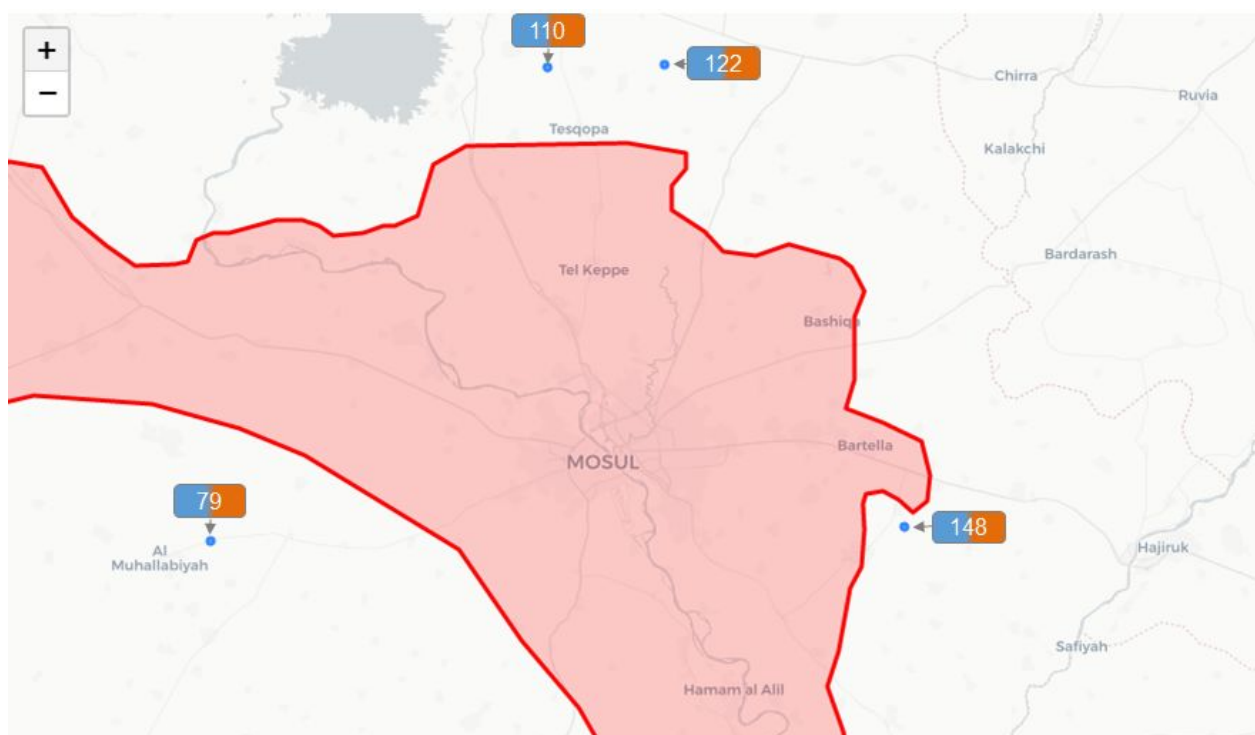
**Figure 5.4:** Most frequently returned individuals after 300 runs evaluating centrality and total population (blue) and 300 runs evaluating centrality and IDPs outside camps (orange).



**Figure 5.5:** Normalised scores of the most frequently returned candidates after 300 runs evaluating centrality and total population (left) and 300 runs evaluating centrality and IDPs outside camps (right).

The impact of the restricted area for the 21 October 2016 timestep is immediately visible by comparing Figures 5.3 and 5.6. The restricted area forced the algorithm to posit locations further away from Mosul, and of the most frequently returned candidates for both evaluation settings, two are different to those from the benchmarking exercise. As will be seen in the next subsections, the restricted area grows quite significantly in subsequent time steps, engulfing the locations of candidates 79 and 148. In practice these would not have been secure locations and reveal some of the limitations of taking the algorithm's results at face value. This issue could be addressed by an astute user of the SDSS in two ways: liaising with coalition forces and taking into account their plan to enter Mosul from the east may facilitate the creation of a revised anticipated restricted area that would exclude candidate 79 from consideration. Candidate 148 is only 1.4km away from the restricted area so using the option to input a minimum distance from this zone - which is available in the SDSS - would help mitigate selection of a potentially risky location.

In terms of the algorithm's behaviour, the most commonly returned individual differed between the two evaluation criteria. For the total population setting, it consisted of candidates 110 and 148, which was delivered in 33% of the 300 runs. 148 was paired with 122 under the IDPs setting, which returned it in 41% of runs. This was because 122 had the maximum normalised IDPs score but the second highest normalised population score (beaten by 110). For these runs the algorithm was working as intended, finding the highest scoring candidate available.




**Figure 5.6:** Visualisation of most frequently returned candidates. Both the total population and IDPs outside camps evaluation settings selected the same locations.

**Table 5.2:** Details of the most frequently returned candidates for 21 October 2016.

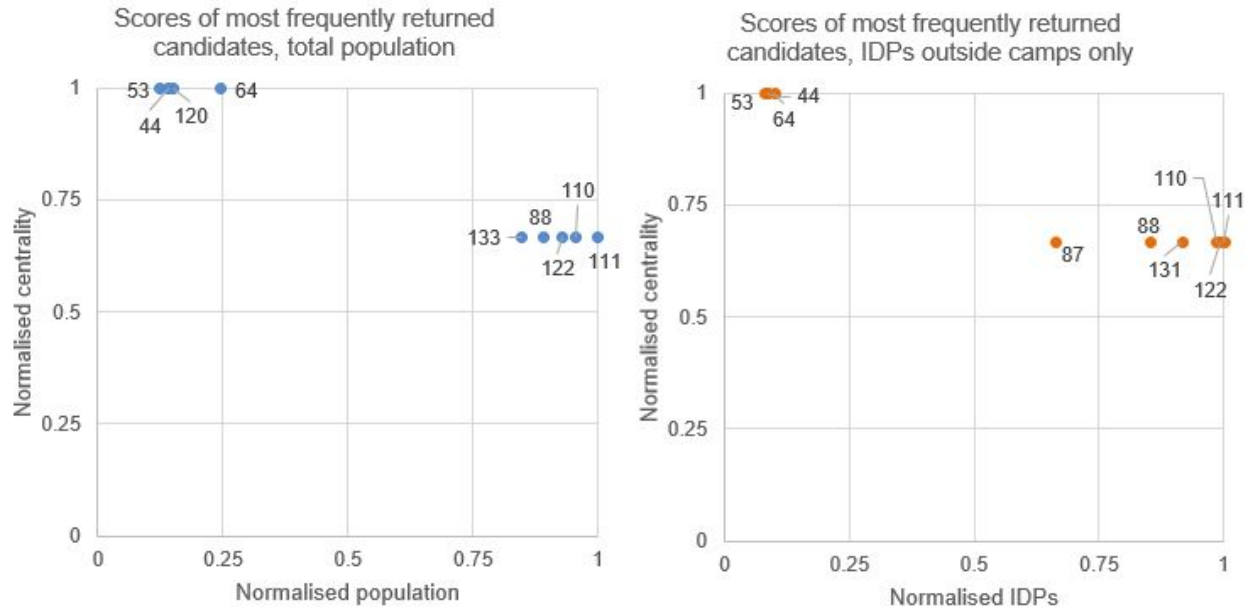
Candidate ID	Total population	IDPs outside camps	Degree centrality	TSPs	Proximity to restricted (km)
110	2,691,844	407,574	3	3	6.48
148	2,366,775	172,824	4	2	1.40
122	2,614,827	410,868	3	2	7.09
79	2,144,843	102,066	4	0	9.84

### 5.1.3 3 November 2016



**Figure 5.7:** Most frequently returned individuals after 300 runs evaluating centrality and total population (blue) and 300 runs evaluating centrality and IDPs outside camps (orange).

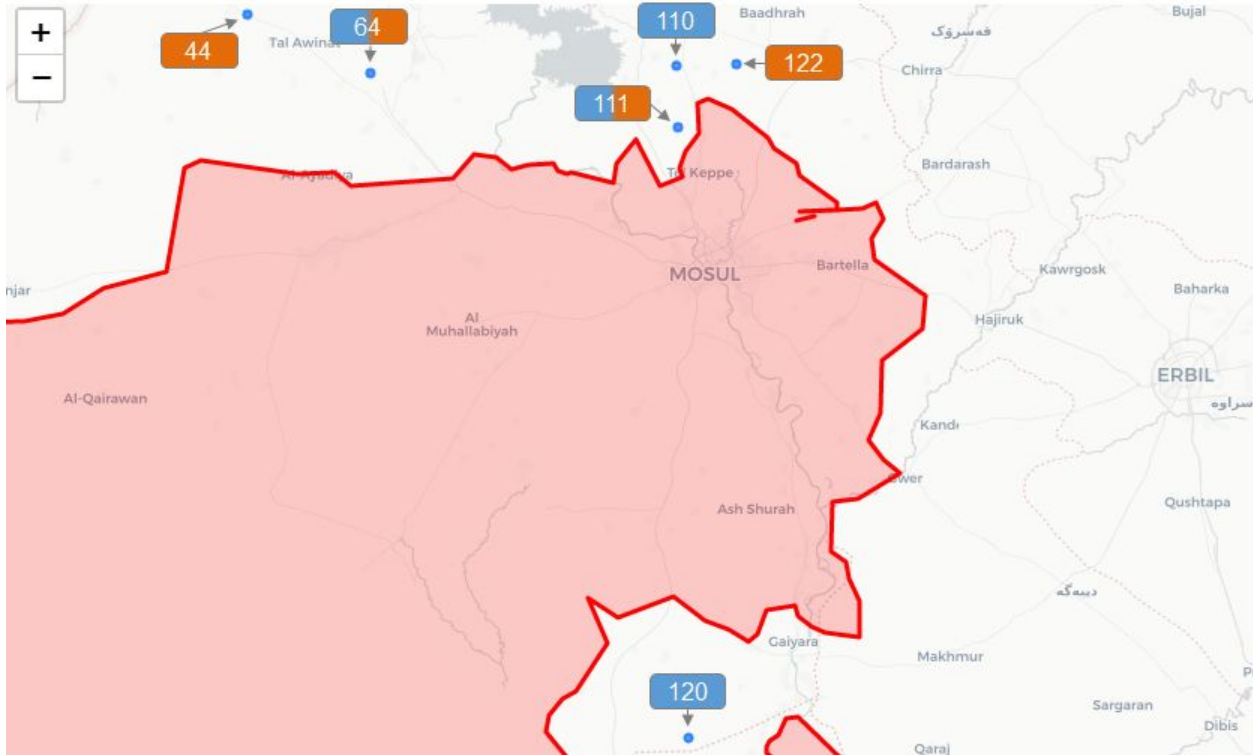




**Figure 5.8:** Normalised scores of the most frequently returned candidates after 300 runs evaluating centrality and total population (left) and 300 runs evaluating centrality and IDPs outside camps (right).

The 3 November 2016 timestep saw the restricted area reach its largest spatial extent. As shown in Figure 5.7, an individual comprised of candidates 111 and 64 was most frequently returned under both the total population and IDPs settings. In comparison to runs for other timesteps, this was the most dominant solution, being delivered 63% of times for the total population setting and 62% for the IDPs setting. It can be inferred that these relatively high percentages are a function of the lower number of candidates available for selection in the starting population. Due to the extent of the restricted areas, this pool was limited to 86 candidates. The visualisation in Figure 5.9 shows how sites to the east of the Mosul that were prevalent among solutions returned for other time periods are unavailable because of the size of the restricted area, which was at its maximum extent for this time step. As shown in Table 5.3, all candidates are within one hour's drive of at least one TSP with the exception of candidate 120. The possible selection of candidate 111 as the optimum site under both evaluation settings - substantiated by its proximity to 4 TSPs and population and IDP count - could be refuted by how close it is to the restricted area. Although it does not fully encroach on this site in subsequent time steps, this is another instance where the user inputting a minimum threshold for distance to the restricted area would have a significant influence upon the results.



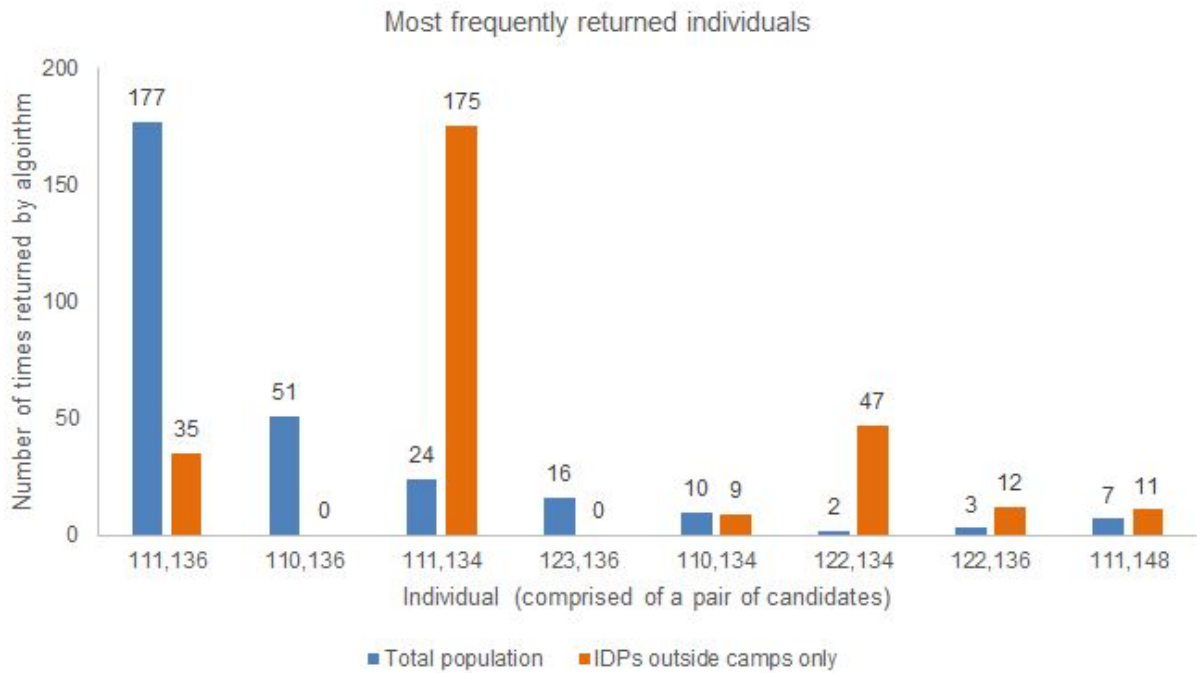


**Figure 5.9:** Visualisation of most frequently returned candidates for runs evaluating total population (blue labels) and IDPs outside camps only (orange labels). Half blue, half orange labels correspond with locations returned using both settings.

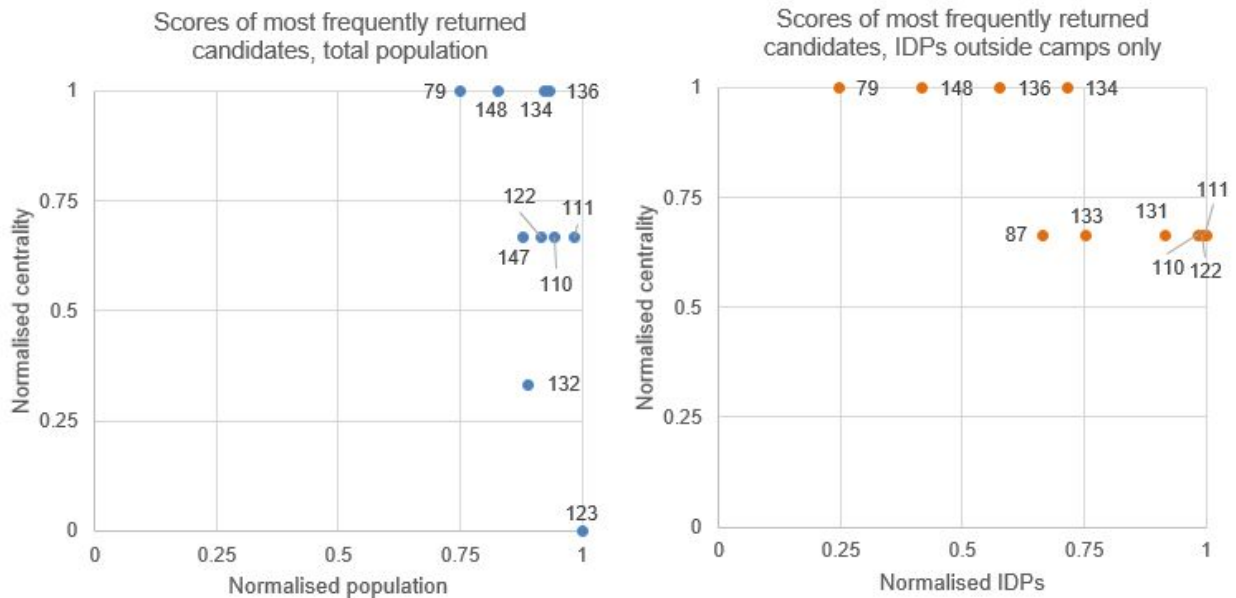
**Table 5.3:** Details of the most frequently returned candidates for 3 November 2016.

Candidate ID	Total population	IDPs outside camps	Degree centrality	TSPs	Proximity to restricted (km)
111	2,805,417	413,856	3	4	3.38
64	792,280	41,910	4	1	17.29
120	538,058	2,688	4	0	12.93
110	2,691,844	407,574	3	3	6.98
122	2,614,827	410,868	3	2	6.97
44	518,087	36,882	4	0	24.51

### 5.1.4 24 November 2016



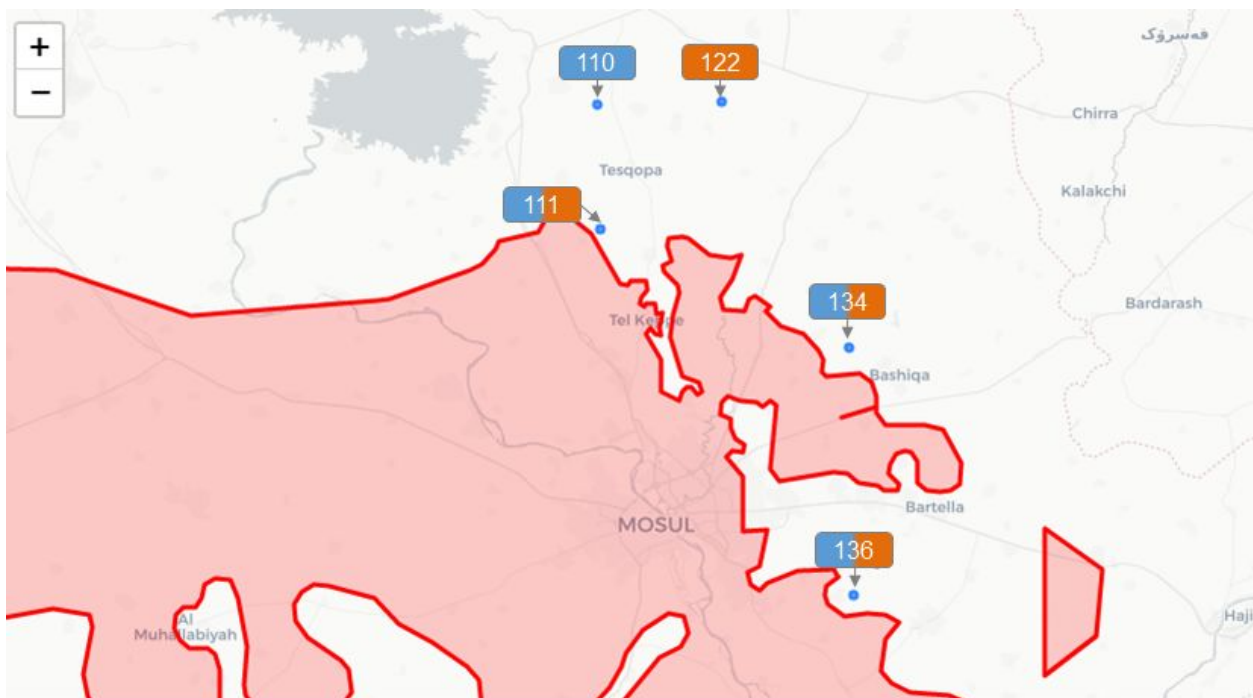
**Figure 5.10:** Most frequently returned individuals after 300 runs evaluating centrality and total population (blue) and 300 runs evaluating centrality and IDPs outside camps (orange).



**Figure 5.11:** Normalised scores of the most frequently returned candidates after 300 runs evaluating centrality and total population (left) and 300 runs evaluating centrality and IDPs outside camps (right).

The 24 November 2016 timestep marked the retreat of ISIS after the initial phases of the coalition assault on Mosul. Referring to Figures 5.10 and 5.1, the top performing individuals under both evaluation settings were identical to those of the benchmarking exercise. However, similarities in the favoured solutions do not tell the whole story: the top solutions for the population setting and the IDPs setting were returned by 59% and 58% of runs compared to 27% and 31% in the benchmarking exercise. This is to be expected since fewer candidates were available for selection as part of the starting population for this timestep than the benchmarking exercise, but it is interesting that the difference in returned frequency of top individuals is so substantial.

Compared with the previous timestep, the re-availability of high performing candidates to the east of Mosul (134 and 136) was evident because the four most frequently returned individuals for both evaluation settings featured one of these candidates. Proximity to the restricted area remained an issue - Table 5.4 shows that both candidates were within 1.61km of the boundary. In addition, Figure 5.12 illustrates that 136 is sited in a location between restricted areas influenced by ISIS, making it vulnerable to encirclement.




**Figure 5.12:** Visualisation of most frequently returned candidates for runs evaluating total population (blue labels) and IDPs outside camps only (orange labels). Half blue, half orange labels correspond with locations returned using both settings.


**Table 5.4:** Details of the most frequently returned candidates for 24 November 2016.

Candidate ID	Total population	IDPs outside camps	Degree centrality	TSPs	Proximity to restricted (km)
136	2,659,086	238,782	4	8	1.61
111	2,805,417	413,856	5	4	0.84
110	2,691,844	407,574	6	3	9.27
134	2,634,518	295,614	7	7	2.08
122	2,614,827	410,868	3	2	11.24

### 5.1.5 9 January 2017



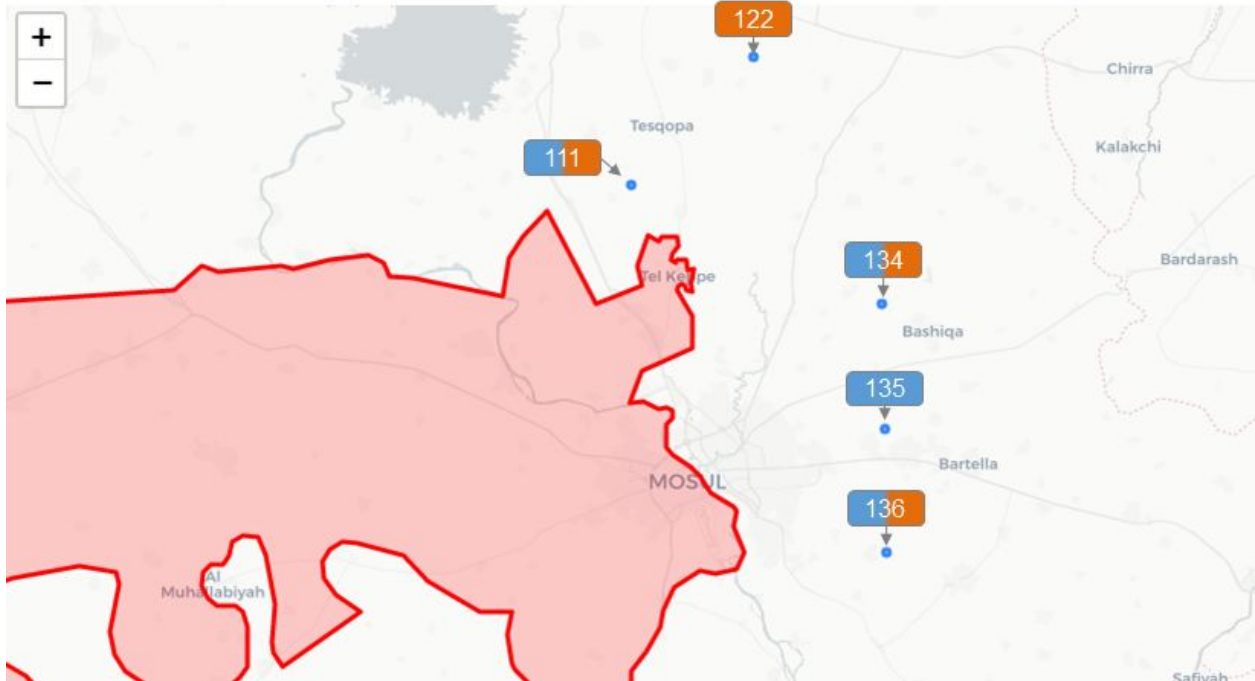
**Figure 5.13:** Most frequently returned individuals after 300 runs evaluating centrality and total population (blue) and 300 runs evaluating centrality and IDPs outside camps (orange).



**Figure 5.14:** Normalised scores of the most frequently returned candidates after 300 runs evaluating centrality and total population (left) and 300 runs evaluating centrality and IDPs outside camps (right).

ISIS continued to retreat from the east of Mosul between November 2016 and January 2017, making the restricted area for 9 January 2017 even smaller. The individuals selected most frequently under the two evaluation settings were the same as in the previous timestep. Candidate 135 was now available and was the fourth most frequently selected candidate overall for the total population setting due to its proximity to Mosul, which under the model was still represented as being densely populated. Conversely, it did not appear at all in the top five individuals for the IDPs setting because its drive time polygon did not encompass the IDPs to the north of Mosul (reachable from 111, 122 and 134) or to Erbil in the east (reachable from 136).

Previous concerns about proximity to the restricted area are less acute for this timestep. Although this timestep marks the final temporal period for the restricted areas data provided by the WHO, it should be pointed out that it preceded the third phase of the battle to liberate Mosul, which saw the highest number of traumatic injuries: 85% of coalition casualties were suffered during this phase (Arnold and Fiore, 2019). Coalition forces were aware when planning the assault that this phase was likely to see the most severe fighting, reinforcing the argument that healthcare providers should favour candidate locations that have the ability to service the most TSPs. If the model was used in a context after fighting had ceased, a strong case can be made for population being the deciding factor between candidate locations.



**Figure 5.15:** Visualisation of most frequently returned candidates for runs evaluating total population (blue labels) and IDPs outside camps only (orange labels). Half blue, half orange labels correspond with locations returned using both settings.

**Table 5.5:** Details of the most frequently returned candidates for 9 January 2017.

Candidate ID	Total population	IDPs outside camps	Degree centrality	TSPs	Proximity to restricted (km)
111	2,805,417	413,856	4	4	4.29
136	2,659,086	238,782	3	8	11.47
134	2,634,518	295,614	3	7	15.29
135	2,703,547	267,882	4	8	13.66
122	2,614,827	410,868	3	2	15.69

## **5.2 Comparison with previous models**

The previous section indicated that the SDSS was capable of returning viable indicative locations for mobile hospitals, with the caveat that users need to use their own judgement regarding proximity to restricted areas. To further consider the usefulness of the system, its results and speed will now be compared with two alternative models that were developed using the same case study. This will be preceded by a brief summary of the approaches taken by these models.

### **5.2.1 Model developed by the ITC / WHO (Amer, Augustijn, van den Bosch and Da Silva Mano (2017))**

The model developed by the ITC at the University of Twente sought to deliver the best possible health benefits to the largest number of IDPs affected by the Mosul battle through provision of secondary health care. This took into account the spatial coverage and numbers of IDPs serviced by existing secondary health facilities, the referral process sending IDPs to facilities depending on their needs, and a location-allocation methodology able to select priority locations for field hospitals. Scenarios were modelled to reflect dynamic factors like IDP movements and territorial control by ISIS forces. The first scenario considered healthcare access before significant IDP movements, the second and third posited situations with 250,000 and 500,000 IDPs respectively (and the influence of optimally located field hospitals), while the fourth looked at a situation where east Mosul had been liberated.

A custom model designed in ArcGIS model builder called CandidateHospitals was used to find the best field hospital locations and the scenarios each used a modified version of this model. The authors acknowledged that the user requires some prior conceptual and practical GIS skills to operate the models. They also recognised the amount of time required to organise data for input into the models, a point pertinent for all three projects being considered here.

### **5.2.2 Model developed by van der Caaij (2019)**

Van der Caaij (2019) took a more interactive approach, developing a SDSS in QGIS that enabled users to manipulate myriad variables influencing field hospital location and service population. He made the distinction between static and dynamic indicators in the model. The former included connectivity and resilience, generated using a custom betweenness centrality measure and a custom degree centrality measure respectively. Dynamic indicators included population, distance to potential danger areas, access to water and space available. Users were also able to manipulate the road network by disabling segments, choose different weightings for suitability and load predefined scenarios.

Greater emphasis was placed on the usability of the model than the ITC / WHO concept, with a dedicated dialogue box to guide the user through the different choices available. Of the three

models considered it was the only one that was formally tested by users representing different stakeholders.

### 5.2.3 Theoretical comparison

Table 5.6 provides a side-by-side overview of the different characteristics of all three models together with some identified strengths and weaknesses. While all three intend to provide a similar output (mobile or field hospital locations) and all had access to the same data, it is evident that each project chose to prioritise particular aspects of the model. For instance, van der Caaij opted to focus on usability, investing time in user testing and development of a smooth interface, whereas this was not the emphasis of the other two models. As stated in the research objective, this project intended to deliver a model able to return quality results faster than the ITC / WHO option. To measure whether or not this has been achieved, the model will be set up with the same data inputs where possible and the results of the two will be compared.

**Table 5.6:** Comparison between the key components of the three models along with their strengths and weaknesses.

<b>Component</b>	<b>Genetic algorithm model (this thesis)</b>	<b>ITC / WHO model</b>	<b>van der Caaij model</b>
Service area of facility	Default of 60km (2 hours drive time at 30km/hour).	50km.	Voronoi diagram to divide map between different hospitals.
Service population	Host population, IDPs not in camps, IDPs in camps or just IDPs not in camps depending on setup.	Host population, IDPs not in camps, IDPs in camps depending on scenario.	Host population, IDPs not in camps, IDPs in camps.
Road network	ITC / WHO network used in generation of candidate grid, new layer downloaded from OpenStreetMap for the model.	OpenStreetMap but additional work carried out to validate network.	OpenStreetMap but additional work carried out to validate network.
Dynamic factors	Changing restricted areas, input minimum threshold to restricted areas.	Four scenarios reflecting different stages of the battle. Changing restricted areas and population location (IDPs and host).	Changing restricted areas, hospital service areas and capacities, suitability weights, make road segments inactive, add own hospital locations.



Main tool for selecting best location(s)	Genetic algorithm.	ArcGIS location allocation tool.	Custom multi-criteria decision making tool.
Interface	Dialogue box for choosing time period and running model. Output in Folium map.	ArcGIS model builder.	Custom application for QGIS including dedicated dialogue components and opportunity to choose weights and thresholds.
Strengths	<p>Lightweight as not reliant on desktop GIS software, fast, easy to switch between timesteps/scenarios, inputs can be substituted if new data is available.</p> <p>High potential for expansion / modification - very straightforward to add other input options such as custom drive time service areas.</p>	Very comprehensive, takes into account multiple variables, good justification of field hospital locations.	Able to add own input layers (including hospital locations) and analyse their effects on hospital performance. Lots of interactive variables. Feedback from real users.
Weaknesses	User requires some GIS and Python knowledge if they want to change inputs significantly. Overlapping service areas if multiple locations selected. Does not take into account existing healthcare network aside from TSPs.	User requires prior GIS knowledge. Scenarios quite rigid.	Concerns from users about speed. Perhaps overly complex. Users require prior GIS knowledge for operation. Voronoi diagram overly simplistic for service areas.

**5.2.4 Setup and test description**

To compare results delivered by the genetic algorithm with the ITC / WHO model the inputs needed to be as close as possible for a fair test. Scenario 0 from the ITC / WHO was chosen to compare against because of the four available scenarios this had the inputs nearest to those of the GA. Scenario 0 was characterised by the following inputs available for use in the GA model:

- Existing host population

- IDPs in camps and outside camps
- All restricted areas dated 3/11/2016

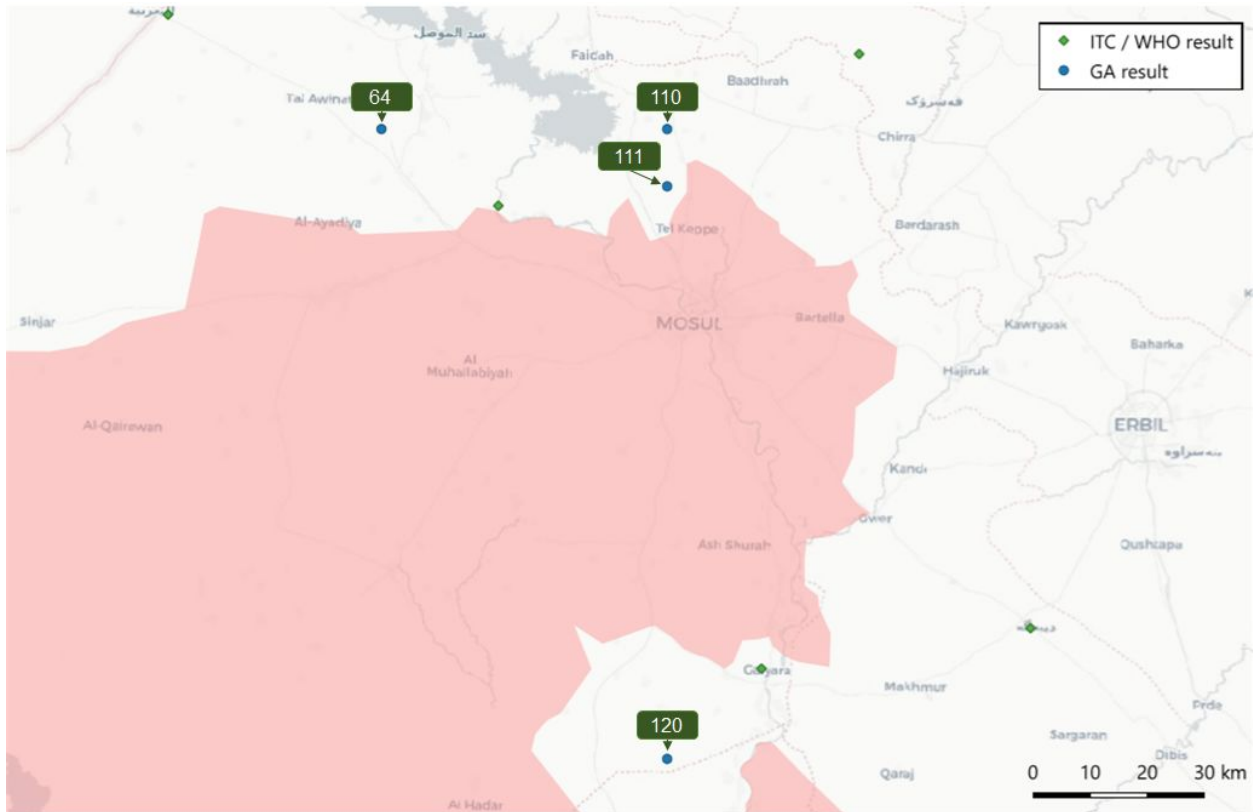
For this scenario, the ITC / WHO model incorporated existing health facilities and barriers to the road network, two inputs that were not available for the genetic algorithm model. Another variation was the calculation of accessibility: the ITC / WHO model also used the road network, but did not take any measure of centrality into account.

Acknowledging these differences, the genetic algorithm model was run 300 times using the following parameters intended to represent Scenario 0 as closely as possible:

- Only candidates outside restricted areas dated 3 November 2016 were evaluated.
- As in section 5.1, runs used a starting population of 80 individuals, each consisting of a random pair of candidates. Each run terminated after 30 generations.
- Normalised degree centrality and normalised population (existing host population, IDPs in camps and outside camps) were evaluated by the algorithm.

Carrying out multiple runs meant that the best individual(s) would be clear to identify and that there would be a good sample for assessing the time taken to return a result.

## 5.2.5 Test results



**Figure 5.16:** Comparison of ITC / WHO Scenario 0 result with the equivalent for the genetic algorithm model.

Figure 5.16 visualises the four most frequently returned candidates alongside the five locations recommended by the ITC / WHO analysis. 63% of the genetic algorithm's 300 runs chose candidates 64 and 111 as the best individual, with this followed by a pairing of 120 and 111 (15% of solutions). Although there are some spatial commonalities between the output of the two models - both favoured locations north of Mosul, for instance - the shortest distance between a genetic algorithm location and an ITC / WHO option is 23km, an amount that could not be considered adjacent considering the size of the study area.

To enable further direct comparison, the five ITC / WHO locations were assessed using the genetic algorithm processes and criteria, with the results presented in Table 5.7. If all locations are taken individually, it is clear that two of the genetic algorithm candidates (111 and 110) were strongest in terms of population served and number of TSPs in the vicinity. However, if the two sets of candidates were taken as different solutions for siting multiple hospitals, the ITC / WHO method might be deemed preferable since it seeks to maximise the cumulative population served by all sites, using distinct service areas to achieve this. The genetic algorithm does not include such a mechanism; in practical terms most of the individuals served by a hospital at

candidate 110 would also be covered by candidate 111 while those elsewhere in the study area might be neglected.

**Table 5.7:** Details of the genetic algorithm’s four most frequently returned candidates for inputs emulating Scenario 0 and the equivalent for the ITC / WHO recommendations (in blue).


Candidate ID / name	Proximity to restricted (km)	Total population	Degree centrality	TSPs	Nearest ITC / WHO result (km)
111	3.38	2,805,417	3	4	29.73
64	17.29	792,280	4	1	24.43
120	12.93	538,058	4	0	22.81
110	6.98	2,691,844	3	3	32.43
Al Qayyarah	4.25	850,565	3	0	N/A
Shekhan	31.31	1,091,017	3	0	N/A
Rabea' Area	34.71	313,814	3	0	N/A
Wana	0.94	1,790,145	3	0	N/A
Dibaga	34.46	682,056	4	0	N/A

### 5.2.6 Measuring time

This project stated in its objective that the SDSS should “rapidly” calculate and display indicative mobile hospital locations, a criteria that led to the selection of a genetic algorithm to underpin the model. Testing the model’s speed and comparing it with previous work enables measurement of the extent that this aspect of the objective has been achieved. However, differences in process order and data preparation between the genetic algorithm model and the ITC / WHO model make a direct comparison difficult to organise in a fair manner. With this in mind, it is important to be explicit about exactly which components of the models were run to collect time data. The processes assessed should not be treated as like-for-like, but the following results are still useful if these caveats are recognised.

Among the various tools developed by the ITC / WHO to reflect different scenarios in the CHE, the model named CandidateHospitals was closest to the genetic algorithm model in terms of its inputs and outputs. This model was constructed in ArcGIS model builder and used the ArcGIS location allocation tool to return hospital locations based on the restricted area, the road


infrastructure, population, IDP camp locations and pre-existing healthcare facilities. The model also generates random points for potential candidate locations, a procedure that takes place in the data preparation phase for the genetic algorithm model. The data supplied with the model covered a larger spatial extent than that of the genetic algorithm model, so inputs were clipped beforehand to cover the same area. As in the genetic algorithm model, the user was able to change the time period to manipulate the restricted area. To gather speed data, the model was run 60 times in ArcGIS model builder: 30 iterations where the restricted area was at its greatest extent (3 November 2016) and 30 where it had shrunk (9 January 2017). Figure 5.17 shows the results of these runs. The model took an average of 28.95 seconds to return a solution for the 3 November setting and an average of 30.9 seconds to complete a run for the 9 January setting.



**Figure 5.17:** Time taken for the ITC / WHO CandidateLocation model to deliver solutions under two different restricted area settings.

To collect speed data on the genetic algorithm model, measurements were taken for the cumulative time it took to conduct the two key processes that make up the live SDSS: the safety and genetic algorithm phases. A Python loop was set up so time data could be gathered for 300 runs for the 3 November timestep and 300 runs for the 9 January timestep. Figure 5.18 displays the results. For the 3 November parameter a solution was returned in an average of 3.61 seconds while an average of 4.06 seconds was needed to deliver a solution under the 9 January setting. These results imply that it can be cautiously suggested that the GA model is

significantly faster than the ITC / WHO CandidateLocation model, with the caveat that there were some fundamental differences in the processes they carry out. These differences and their implications will be expanded upon in the discussion section.



**Figure 5.18:** Time taken for the GA model to deliver solutions under two different restricted area settings.

This chapter collected data and offered commentary demonstrating that different timesteps have a big influence on the model's results (sub-chapter 5.1) and presented a comparison between the model and previous work (sub-chapter 5.2). Drawing on this evidence, the following chapter will consider the meaning and significance of these results.

## 6. Discussion

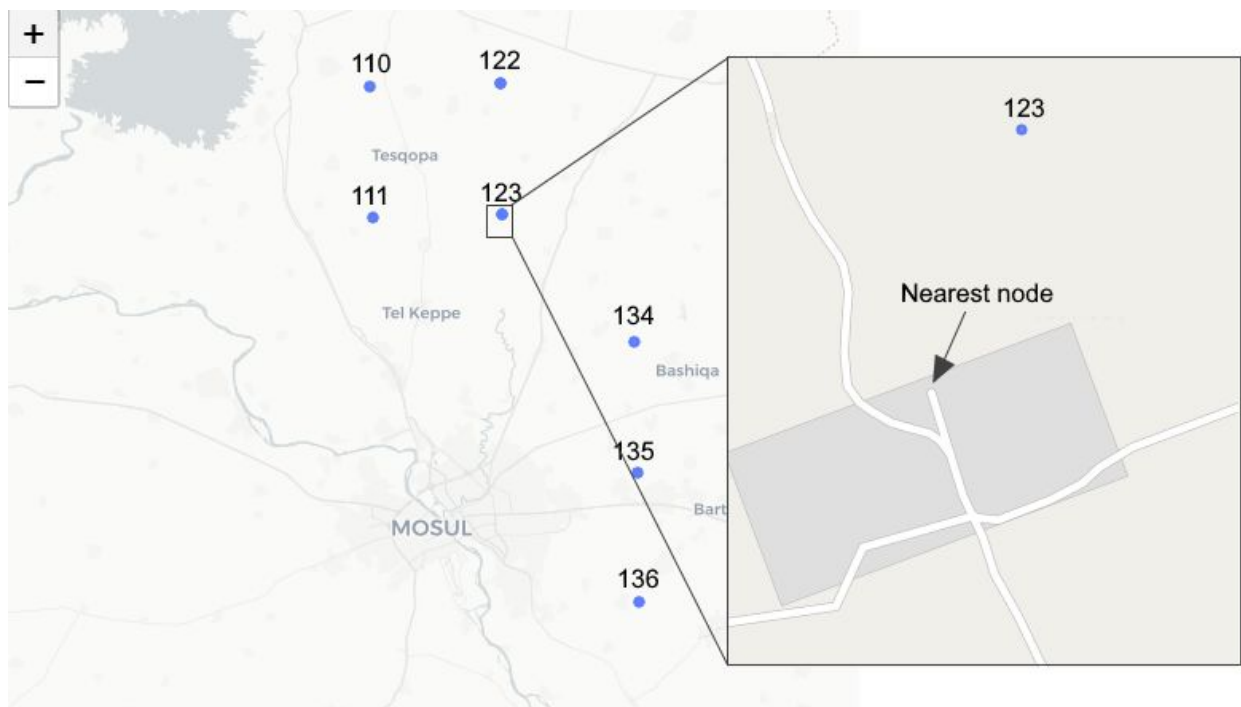
### 6.1 Quality and temporal variation

The results presented in sub-chapter 5.1 provided good evidence that the model was capable of consistently returning solutions to the location allocation problem. The heuristic nature of the underlying genetic algorithm meant that there was some variation in solutions delivered over the 300 runs for each timestep. The benchmarking exercise saw the biggest spread of solutions with five individuals returned by 7% or more of the 300 total population runs while the best solution achieved a 31% share. The best solutions delivered for each of the four timesteps were noticeably more prevalent, ranging from a 33% share (21 October) to 63% share (3 November). The runs evaluating IDPs outside camps exhibited similar behaviour. For both settings, inverse correlation between dominance of the solution and the amount of candidates available for selection in the starting population can be observed: for the benchmarking exercise, all 183 candidates could be chosen, while the restricted area was at its greatest extent by 3 November, meaning only 86 candidates were available. This would become particularly important if the model were applied to a new case study with a bigger candidate pool. This also indicates the usefulness of the initial filtration stage of the candidate location grid using the results of the supplied candidates analysis. Inclusion of all 252 candidates from the original grid would further disrupt the frequency the algorithm arrived at a strong solution.

Taking a closer look at the formulation of the individuals proposed as best solutions for each timestep further reveals the extent that recommended locations changed as the CHE unfolded. Under the total population criteria of the model - safety as determined by restricted areas, demand from population and centrality of the nearest node - there was no single location that featured in the top two individuals across all timesteps. For the IDPs setting, the only location that appeared in the top two individuals for all four timesteps and the baseline exercise was candidate 122. This was because this site never fell in the restricted area during the CHE. For both settings, the differences between the top four locations suggested on 21 October and 3 November are particularly apparent. While the earlier timestep proposed candidates to the north, east and west of Mosul city, the growth of the restricted area by 3 November meant that sites to the east and west were no longer viable. Perhaps this is an obvious point but it stresses the challenging nature of location allocation during CHEs - within a matter of weeks, movements by hostile forces (or fresh intelligence on the ground situation) compromised locations previously perceived as optimal. Making predictions about how the restricted area (or indeed dispersion of population or IDPs) change over time is clearly beyond the model's scope, but these changes do reinforce the importance of the model's speed and ability to incorporate new data sets. As will be explained during discussion of the model comparison results, it was largely successful at doing this.

Having acknowledged that the model is capable of reaching a solution and that these vary between timesteps, it is worth focusing on the practical viability of these suggested locations.

The model is necessarily an abstraction that only incorporates some of the identified needs of a mobile hospital. The decision to weight centrality and population equally within the algorithm meant that candidate 123, which of all candidates had the highest population and IDPs within its service area, never appeared within the top solutions for any of the timestep scenarios because it is located at a deadend, so scored poorly for centrality. Figure 6.1 shows that in reality it is also close to a crossroads that achieved the maximum centrality score. Consequently, a valid case could be made for either modifying the weighting of centrality or devising an improved measurement for this metric. Van der Caaij (2019) also recognised the usefulness of degree centrality but extended it to provide an indication of resilience by taking an average of all nodes within 200m of the candidate location. Experimentation with similar approaches could lead to locations like candidate 123 receiving a higher centrality score.



**Figure 6.1:** Candidate 123 in relation to Mosul and its nearest node. Note how its position on the grid is adjacent to candidates frequently returned as a solution (122 and 111).

Regarding the population metric, the data collection for the results section either used IDPs not in camps or total population (host population plus IDPs) as the population inputs. The IDPs not in camps group was singled out because of their vulnerability to healthcare issues but this also demonstrated the model's ability to ingest alternative datasets for the population variable. What is worth pointing out is the difficulty of accounting for population movements between timesteps. Previous work shows that a significant proportion of Mosul's population were displaced during the ISIS regime and especially the CHE resulting from the struggle to liberate the city, but the available spatial data for both the host population and IDPs did not include temporal changes. Amer et al. (2017) addressed this by modelling their own scenarios; this model was also able to accommodate alternative population data layers. However, simulating population movements to



create these layers is a demanding undertaking in itself. The lesson here - backed up by the conclusions of Amer et al. (2017) - is that the model's recommendations are directly tied to the quality of the data inputted, and in a fast-moving CHE this cannot be guaranteed. Indeed, the amount of data available for the case study was surprising and possibly exceptional for the circumstances. Users seeking to deploy the model in future CHEs may find sourcing data significantly more challenging.

## **6.2 Comparison with previous models**

Comparison with previous work was a useful way of exploring the strengths and weaknesses of the genetic algorithm model. It factored in fewer layers than the models devised by Amer et al (2017) and van der Caaij (2019) but appeared to be faster and more lightweight. This can be attributed to the use of a genetic algorithm rather than more conventional overlay techniques and the decision to construct it purely in Python instead of relying on a desktop GIS. The longest run in the speed tests took 6.92 seconds but the average time was 3.84 seconds to complete a run for the two evaluated timesteps. The equivalent measurements for the ITC / WHO model were 36.27 and 29.92 seconds respectively. Although the differences in processes between the models should again be stressed, this speed disparity remains notable. For stakeholders in a CHE using an SDSS to locate mobile hospitals, who need to quickly home in on a solution and potentially want to test a variety of data inputs, the genetic algorithm model would clearly be preferable.

Systematic usability testing was out of scope for this project but van der Caaij (2019) demonstrated that there was a lot to be learned from a focus group approach, especially given the range of stakeholders concerned with this location allocation problem. In terms of output for Scenario 0, under the criteria of this project candidates 111 and 110 were better individual locations than some of those selected by the ITC / WHO because they had bigger populations in their service area. However, as a collective solution there were concerns about overlapping service areas. The ITC / WHO method delineated distinct service areas for each site and as a result provided better healthcare coverage of the study area. It can be inferred from this comparison that one of the SDSS' main weaknesses is that it does not have a mechanism that recognises the spatial relationship between the service areas of candidate locations. While the best performing candidates were still viable, a user seeking to position multiple hospitals would need to perform additional analysis to determine the best combination of candidates.

## **7 Conclusion, limitations and opportunities for future work**

### **7.1 Conclusions**

This research sought to develop an SDSS that enabled the user to quickly display indicative suitable locations for siting one or more mobile hospitals that meet the needs of the greatest number of patients during and after a CHE, focusing on the case study of the Mosul liberation. The results and discussion suggest that it was largely successful at realising this and adds weight to the argument in favour of using genetic algorithms to solve location allocation problems, especially when speed is a priority. Evidence for this claim can be found by returning to the research questions and summarising the answers that were arrived on during the project.

#### **7.1.1 What are the requirements of a mobile hospital and in which order should these be prioritised?**

This question was addressed during the literature review and the analysis of the 20 locations supplied by the WHO. Mobile hospitals have myriad requirements which broadly fall under the categories of demand, how they relate to the existing healthcare network, site accessibility, site safety, connection to utility networks and the type of space available. Given the central objective of creating a fast model it was necessary to focus on the most important of these so the number of required data inputs and processing time could be minimised. Filtering candidates using parameters taken from the analysis of the supplied locations ensured that from the outset all prospective sites met basic criteria such as proximity to a road. The literature implied that demand, accessibility and safety were the three requirements of particular significance. Safety was highlighted as the most important overall because of the volatile nature of the CHE environment. When designing the SDSS it was deliberately given its own phase in the evaluation process to ensure that any risky locations were discarded immediately.

#### **7.1.2 How does a mobile hospital relate to the wider healthcare network, including Trauma Stabilisation Points and established hospitals?**

The literature review also helped answer this question. For mobile hospitals to have a high efficacy at alleviating public health issues during a CHE, they need to operate in conjunction with what remains of the existing healthcare network. In a CHE their capacity for trauma treatment is often leveraged making short travel time from TSPs an important consideration. The number of TSPs within the 'golden hour' of each candidate location was calculated to gauge this. However, this information was not used in the formal evaluation process, featuring instead as attribute data in the results map. An improved version of the model might integrate this properly, perhaps by filtering the initial candidate pool to those in reach of a TSP.

### **7.1.3 Which methods can be used to solve location allocation problems quickly and effectively, factoring in dynamic variables as the CHE changes over time?**

A process for generating candidate locations, derived from an analysis of supplied locations, was married up with a genetic algorithm as a method able to bypass some of the time consuming aspects of conventional location allocation models. Of available genetic algorithms, NSGA-II's ability to evaluate multiple criteria meant it was especially prominent in previous work and made it an appropriate choice for this project. The decision to pursue this option led to Python being chosen as the language for creating the SDSS which was also advantageous in terms of speed and flexibility. Although some pre-processing in QGIS was necessary for preparing data inputs, the resulting SDSS was capable of calculating results for different temporal scenarios.

### **7.1.4 To what extent do the results returned by the algorithm change as the CHE unfolds?**

The main temporal variable available in the data provided was the location and shape of restricted areas. Tests conducted for the first sub-chapter of the results section demonstrated that these had a profound impact on the solutions returned by the algorithm, with some locations becoming unviable as territorial boundaries shifted over time. This highlighted the difficulty of the original location allocation problem and the necessity for decision makers to have access to an SDSS that is able to be quickly updated with new restricted areas. It is likely that data on population movements during the CHE would cause the results to change even more dramatically, but temporal variation for this data input was not a feature of the original data package.

### **7.1.5 How can the algorithm be integrated with an SDSS to facilitate decision making for non-experts?**

One of the potential issues with electing to write the algorithm in Python rather than using a GIS as a foundation was the probable lack of familiarity non-experts have with Python code and the absence of an intuitive interface. To mitigate this, a basic GUI was wrapped around the algorithm so that users could choose the time period, minimum distance to the restricted area, run the algorithm and see a map showing the results. As a whole, the algorithm, GUI and output window comprised the SDSS, although it should be noted that some setup (installation of certain Python libraries and connection to a PostGIS database) was needed for it to work on another computer. The simplicity of the GUI and the clarity of the results output suggest that the SDSS would be able to facilitate decision making for non-experts, but this could be verified and improved with usability testing like that carried out by van der Caaij (2019).

### **7.1.6 How do the results and performance of the algorithm compare to previous work in this field?**

To answer this question, a comparison of inputs, methodology, strengths and weaknesses was carried out with models designed by ITC / WHO (2017) and van der Caaij (2019). This exercise suggested that the SDSS was significantly faster than the previous models model but not as comprehensive when it came to incorporating different variables for location allocation. A comparison of results between the SDSS and the ITC / WHO models also revealed shortcomings with the model's ability to recommend locations for multiple hospitals operating simultaneously. Regarding the van der Caaij model, the drive time service area polygons were an improvement on the voronoi polygons for calculating population since they used the road network rather than euclidean distance. However, the van der Caaij model had greater interactivity - users had a considerable number of variables they could alter, including editing restricted areas - and a slightly more sophisticated procedure for defining centrality. Identification of these issues is helpful for considering how this work can be built upon, which will be the final area of discussion.

## **7.2 Limitations and opportunities for future work**

The design process behind the SDSS set out to create a model that would be applicable to future CHEs, not just the Mosul case study. It can also serve as a foundation for an improved SDSS that addresses some of the limitations identified during the model development and discussion chapters. Most of these limitations can be perceived as opportunities for future work, and in that spirit are consolidated here, along with thoughts on further advances to the academic conversation.

### **7.2.1 Data inputs and functionality**

The most obvious limitation which has already been described extensively is the lack of any mechanism in the model to prevent overlap of service areas when returning multiple locations. There is already a device for penalising individuals comprised of the same candidate locations, so this could be extended to include those with the same service area polygons. This could also be an opportunity to implement a more sophisticated definition of service area that goes beyond the drive time polygons, assessing the proportion of individuals within the service area who are currently underserved by healthcare facilities, for example. In a similar vein, taking a more nuanced approach towards centrality rather than simply scoring a site based on the nearest node could yield better quality results.

One of the strengths of the SDSS was the relative ease that new restricted area, population and TSP data layers can be uploaded and taken into account by the model, with the caveat that the user does need some Python knowledge and that pre-processing may be required for any new layers. A viable improvement would be the addition of an upload function to the GUI, so users could quickly experiment with different scenarios instead of having to store them in PostGIS and

call them using Python. The road network layer, which was downloaded directly from OpenStreetMap, is technically simple to refresh but significantly more time consuming than the other layers due to the amount of information it contains. The ITC / WHO (2017) team also recognised that it contained inaccuracies such as missing edges, and in a CHE it is probable that the accessibility of some routes would change over time due to roadblocks and surface damage. To address these concerns, a solution would be for users to be able to edit the network inside the model by removing or adding edges. This is not a trivial addition and would require some careful thought to be implemented properly.

### **7.2.2 User requirements and practical implementation**

The SDSS largely fulfills the project objective, but additional testing with real users is recommended before deploying it in any new, unfolding CHE. A formal systematic testing process with stakeholders involved in locating mobile hospitals would further understanding of the requirements of the application, check whether the evaluation choices are appropriate, and consider the validity of the recommended locations. Assessment of user interaction with the SDSS would also inform modification of the GUI and workflow if necessary. Even allowing for these additional tests, it should be reiterated that the results delivered by the SDSS are indicative, and any deployment of mobile hospitals should be preceded by an on-ground site assessment of the recommended location.

### **7.2.3 Genetic algorithm**

The project reinforced the notion that genetic algorithms are a fast and effective foundation for an SDSS. However, it is clear that further refinements could be made. A first step here would be testing different crossover and mutation probabilities to see how they influence the speed and results of the SDSS. Larger candidate pools - or even completely different case studies - and stirring operators similar to the one designed by Huang and Wen (2014) might also be fruitful avenues to pursue. Finally, as recognised elsewhere, other criteria could be introduced for evaluation. The algorithm used in this SDSS was the version of NSGA-II that came included in the DEAP library: this was sufficient for the needs of this project, but other types of genetic algorithm could be worth investigating, especially if there are additional evaluation criteria.

## 8. References

Aitken, P., Leggat, P., Robertson, A., Harley, H., Speare, R., & Leclercq, M. (2009). Health and safety aspects of deployment of Australian disaster medical assistance team members: results of a national survey. *Travel medicine and infectious disease*, 7(5), 284-290.

Alp, O., Erkut, E., & Drezner, Z. (2003). An efficient genetic algorithm for the p-median problem. *Annals of Operations research*, 122(1-4), 21-42.

Amer, S., Augustijn, E., van den Bosch, F. and Da Silva Mano, A. (2017). A geospatial scenario based methodology to ensure optimal accessibility and availability of secondary health care for IDPs of the Mosul operation. University of Twente.

Amouri, O. F., & Reed, P. (2018). Notes from a field hospital south of Mosul. *Globalization and health*, 14(1), 27.

Arnold, M. T. D., & Fiore, M. N. (2019). Five Operational Lessons from the Battle for Mosul. *Military Review*.

Attia, A. A., & Horáček, P. (2001, June). Adaptation of genetic algorithms for optimization problem solving. In *7th International Conference on Soft Computing, Mendel* (pp. 36-41).

Bakowski, J. (2016). A mobile hospital—its advantages and functional limitations. *International Journal of Safety and Security Engineering*, 6(4), 746-754.

Balcik, B., & Beamon, B. M. (2008). Facility location in humanitarian relief. *International Journal of Logistics*, 11(2), 101-121.

Barakat, S., & Ellis, S. (1996). Researching under fire: issues for consideration when collecting data and information in war circumstances, with specific reference to relief and reconstruction projects. *Disasters*, 20(2), 149-156.

Barik, R. K., Samaddar, A. B., & Gupta, R. D. (2009, February). Investigations into the efficacy of open source GIS software. In *Map World Forum*.

Beaubien, J. (2018, February 16). Ethical Dilemma Over Treating Civilians Injured In The Battle For Mosul. *Npr.org*. Accessed 17/9/2019 via <https://www.npr.org/sections/goatsandsoda/2018/02/16/586450445/ethical-dilemma-over-treating-civilians-injured-in-the-battle-for-mosul?t=1568708622461>

Blackwell, T., & Bosse, M. (2007). Use of an innovative design mobile hospital in the medical response to Hurricane Katrina. *Annals of emergency medicine*, 49(5), 580-588.

- Boeing, G. (2017). OSMnx: New methods for acquiring, constructing, analyzing, and visualizing complex street networks. *Computers, Environment and Urban Systems*, 65, 126-139.
- Brennan, R. J., & Nandy, R. (2001). Complex humanitarian emergencies: a major global health challenge. *Emergency medicine*, 13(2), 147-156.
- Bricknell, M. C., & MacCormack, T. (2005). Military approach to medical planning in humanitarian operations. *Bmj*, 330, 1437-1439.
- Callimachi, R. and Mills, A. (Presenter and producer). (2018, May 31). Chapter Seven: Mosul. *Caliphate*. [Audio podcast from *The New York Times*]
- Cetinkaya, C., Özceylan, E., Erbaş, M., & Kabak, M. (2016). GIS-based fuzzy MCDA approach for siting refugee camp: A case study for southeastern Turkey. *International Journal of Disaster Risk Reduction*, 18, 218-231.
- Chan, J., & Teknomo, K. (2016). Hub Identification of the Metro Manila Road Network Using PageRank. *arXiv preprint arXiv:1609.01464*.
- Cooper, L. (1963). Location-allocation problems. *Operations research*, 11(3), 331-343.
- Deb, K., Agrawal, S., Pratap, A., & Meyarivan, T. (2000, September). A fast elitist non-dominated sorting genetic algorithm for multi-objective optimization: NSGA-II. In *International conference on parallel problem solving from nature* (pp. 849-858). Springer, Berlin, Heidelberg.
- Deep, K., & Thakur, M. (2007). A new mutation operator for real coded genetic algorithms. *Applied mathematics and Computation*, 193(1), 211-230.
- De Rainville, F. (2018, February 8). Using different scales for fitness values. Deap - users. Accessed 28/1/2020 via <https://groups.google.com/forum/#!topic/deap-users/GDiDzvRd7yc>
- Diaz-Gomez, P. A., & Hougen, D. F. (2007, June). Initial Population for Genetic Algorithms: A Metric Approach. In *Gem* (pp. 43-49).
- Duarte, A., Henriques, R., & Ribeiro, S. (2019). Use of different optimization algorithms to define service areas of police stations in Portugal. *Evidence-based territorial policymaking: formulation, implementation and evaluation of policy*, 108-115.
- Fox, H., Stoddard, A., Harmer, A., & Davidoff, J. (2018). Emergency Trauma Response to the Mosul Offensive, 2016-2017: A Review of Issues and Challenges. *Humanitarian Outcomes*, 10.

Friedrich, C. J., & Weber, A. (1929). *Alfred Weber's theory of the location of industries*. University of Chicago Press.

Gilbert, G., & Rusch, A. M. (2017). Creating safe zones and safe corridors in conflict. *Policy Brief 5*.

Haghani, A. (1996). Capacitated maximum covering location models: Formulations and solution procedures. *Journal of advanced transportation*, 30(3), 101-136.

Hakimi, S. L. (1964). Optimum locations of switching centers and the absolute centers and medians of a graph. *Operations research*, 12(3), 450-459.

Hanlon, M. (2017). Doctors Without Borders' innovative mobile hospital on a trailer. Newsatlas.com. Accessed 30/9/2019 via <https://newatlas.com/medecins-sans-frontieres-mobile-operating-surgical-trailer-must/50242/>

Hinsley, D. E., Rosell, P. A. E., Rowlands, T. K., & Clasper, J. C. (2005). Penetrating missile injuries during asymmetric warfare in the 2003 Gulf conflict. *British Journal of Surgery: Incorporating European Journal of Surgery and Swiss Surgery*, 92(5), 637-642.

Holland, J. H. (1992). *Adaptation in natural and artificial systems: an introductory analysis with applications to biology, control, and artificial intelligence*. MIT press.

Hosage, C. M., & Goodchild, M. F. (1986). Discrete space location-allocation solutions from genetic algorithms. *Annals of Operations Research*, 6(2), 35-46.

Hospitainer. (n.d., a). Products & Services > Products > Wheels. Accessed 27/9/2019 via <https://hospitainer.com/products-and-services/products/wheels/>

Hospitainer. (n.d., b). Projects > Mobile Hospitainers On Their Way For War Victims in Mosul Iraq. Accessed 1/10/2019 via <https://hospitainer.com/projects/5-mobile-hospitainers-on-their-way-for-war-victims-in-mosul-iraq/>

Hospitainer. (n.d., c). UN Level 2 Mobile Hospital. Accessed 30/9/2019 via <https://hospitainer.com/un-level-2-mobile-hospital/>

Houck, C. R., Joines, J. A., & Kay, M. G. (1996). Comparison of genetic algorithms, random restart and two-opt switching for solving large location-allocation problems. *Computers & Operations Research*, 23(6), 587-596.



- Huang, C. Y., & Wen, T. H. (2014). Optimal installation locations for automated external defibrillators in Taipei 7-Eleven stores: using GIS and a genetic algorithm with a new stirring operator. *Computational and mathematical methods in medicine*, 2014.
- ICRC [International Committee of the Red Cross]. (2006). Mobile health units. Methodological approach. Geneva: ICRC.
- IOM [International Organization for Migration]. (2017). Mosul crisis. Population movements analysis. October 2016 to June 2017. Baghdad: IOM.
- Jackson, L. E., Rouskas, G. N., & Stallmann, M. F. (2007). The directional p-median problem: Definition, complexity, and algorithms. *European Journal of Operational Research*, 179(3), 1097-1108.
- Johnson, R. J. (2015). Post-cold war United Nations peacekeeping operations: a review of the case for a hybrid level 2+ medical treatment facility. *Disaster and military medicine*, 1(1), 1-8.
- Kahraman, C., Ruan, D., & Doğan, I. (2003). Fuzzy group decision-making for facility location selection. *Information sciences*, 157, 135-153.
- Kakakhan, J., Hatahit, W., Musani, A., Amer, S., Augustijn, E. W., van den Bosch, F., & Da Silva, M. A. (2018). Geospatial Modeling to plan humanitarian response and ensure access to healthcare: the case of the Mosul liberation operation. In *Practical Geography and XXI Century Challenges* (pp. 250-251).
- Kim, J., & Yoo, S. (2019). Software review: DEAP (Distributed Evolutionary Algorithm in Python) library. *Genetic Programming and Evolvable Machines*, 20(1), 139-142.
- Kreiss, Y., Merin, O., Peleg, K., Levy, G., Vinker, S., Sagi, R., Abargel, A., Bartal, C., Lin, G., Bar, A., Schwaber, M.J., Ash, N. & Bar-On, E. (2010). Early disaster response in Haiti: the Israeli field hospital experience. *Annals of internal medicine*, 153(1), 45-48.
- Kumar, M., Husian, M., Upreti, N., & Gupta, D. (2010). Genetic algorithm: Review and application. *International Journal of Information Technology and Knowledge Management*, 2(2), 451-454.
- Lafta, R., Al-Nuaimi, M. A., & Burnham, G. (2018). Injury and death during the ISIS occupation of Mosul and its liberation: Results from a 40-cluster household survey. *PLoS medicine*, 15(5).
- Larrañaga, P., Poza, M., Yurramendi, Y., Murga, R. H., & Kuijpers, C. M. H. (1996). Structure learning of Bayesian networks by genetic algorithms: A performance analysis of control parameters. *IEEE transactions on pattern analysis and machine intelligence*, 18(9), 912-926.

Leidig, M., & Teeuw, R. (2015). Free software: A review, in the context of disaster management. *International Journal of Applied Earth Observation and Geoinformation*, 42, 49-56.

Li, X., & Yeh, A. G. O. (2005). Integration of genetic algorithms and GIS for optimal location search. *International Journal of Geographical Information Science*, 19(5), 581-60

Manoochchery, S., Hoseinzadeh, E., Taha, P., Rasouli, H. R., & Hoseinzadeh, S. (2018). Field Hospital in Disasters: A Systematic Review. *Trauma Monthly*, 24(2).

MCC [Medical Consultancy & Construction Group]. (n.d.). Mobile Hospitals. Accessed 27/9/2019 via <https://mccgrouplondon.com/project/mobile-hospitals-4/#prettyPhoto>

Mekonnen, A. D., & Gorsevski, P. V. (2015). A web-based participatory GIS (PGIS) for offshore wind farm suitability within Lake Erie, Ohio. *Renewable and Sustainable Energy Reviews*, 41, 162-177.

Michlig, G. J., Lafta, R., Al-Nuaimi, M., & Burnham, G. (2019). Providing healthcare under ISIS: A qualitative analysis of healthcare worker experiences in Mosul, Iraq between June 2014 and June 2017. *Global public health*, 1-14.

Mladenović, N., Brimberg, J., Hansen, P., & Moreno-Pérez, J. A. (2007). The p-median problem: A survey of metaheuristic approaches. *European Journal of Operational Research*, 179(3), 927-939.

Modica, G., Pollino, M., Lanucara, S., La Porta, L., Pellicone, G., Di Fazio, S., & Fichera, C. R. (2016, July). Land suitability evaluation for agro-forestry: definition of a web-based multi-criteria spatial decision support system (MC-SDSS): preliminary results. In *International Conference on Computational Science and Its Applications* (pp. 399-413). Springer, Cham.

OCHA. (2017). Guide for the Military 2.0. *United Nations Humanitarian Civil-Military Coordination*. Geneva, Switzerland: United Nations Civil-Military Coordination Section.

Oppio, A., Buffoli, M., Dell'Ovo, M., & Capolongo, S. (2016). Addressing decisions about new hospitals' siting: a multidimensional evaluation approach. *Annali dell'Istituto superiore di sanita*, 52(1), 78-87.

Pandey, H. M., Chaudhary, A., & Mehrotra, D. (2014). A comparative review of approaches to prevent premature convergence in GA. *Applied Soft Computing*, 24, 1047-1077.

Pérez, C. R. (2015). Emergency response to earthquake in Chile: Experience of a Cuban field hospital. *MEDICC review*, 17, 39-42.

- Piszcz, A., & Soule, T. (2006). Genetic programming: Optimal population sizes for varying complexity problems. In *Proceedings of the 8th annual conference on Genetic and evolutionary computation* (pp. 953-954). ACM.
- Raad, I. I., Chaftari, A. M., Dib, R. W., Graviss, E. A., & Hachem, R. (2018). Emerging outbreaks associated with conflict and failing healthcare systems in the Middle East. *Infection Control & Hospital Epidemiology*, 39(10), 1230-1236.
- Rachmawati, L., & Srinivasan, D. (2009). Multiobjective evolutionary algorithm with controllable focus on the knees of the Pareto front. *IEEE Transactions on Evolutionary Computation*, 13(4), 810-824.
- Rogers, F. B., Rittenhouse, K. J., & Gross, B. W. (2015). The golden hour in trauma: dogma or medical folklore?. *Injury*, 46(4), 525-527.
- Salihy, R. (2019). Terror and Torment: The Civilian Journey to Escape Iraq's War Against the "Islamic State". In Eriksson, J. & Khaleel, A. (Eds.), *Iraq After ISIS* (pp. 79-98). Palgrave Pivot, Cham.
- Santerre, K. D. (1989). From Confiscation to Contingency Contracting: Property Acquisition on or Near the Battlefield. *Military Law Review*, 124, 111.
- Scheffer, D. (2003). Future Implications of the Iraq Conflict: Beyond Occupation Law. *American Journal of International Law*, 97(4), 842 - 860.
- Schutte, S. (2017). Geographic determinants of indiscriminate violence in civil wars. *Conflict management and peace science*, 34(4), 380-405.
- Schwaab, J., Deb, K., Goodman, E., Lautenbach, S., van Strien, M. J., & Grêt-Regamey, A. (2018). Improving the performance of genetic algorithms for land-use allocation problems. *International Journal of Geographical Information Science*, 32(5), 907-930.
- Senvar, O., Otay, I., & Bolturk, E. (2016). Hospital site selection via hesitant fuzzy TOPSIS. *IFAC-PapersOnLine*, 49(12), 1140-1145.
- Serra, D., & Marianov, V. (1998). The p-median problem in a changing network: the case of Barcelona. *Location Science*, 6(1-4), 383-394.
- Shariff, S. R., Moin, N. H., & Omar, M. (2012). Location allocation modeling for healthcare facility planning in Malaysia. *Computers & Industrial Engineering*, 62(4), 1000-1010.
- Sharifi Noorian, S., Psyllidis, A., & Bozzon, A. (2018). A time-varying p-median model for location-allocation analysis. In Mansourian, A., Pilesjö, P., Harrie, L., & von Lammeren, R.

(Eds.). *Geospatial Technologies for All : short papers, posters and poster abstracts of the 21th AGILE Conference on Geographic Information Science. Lund University 12-15 June 2018*, Lund, Sweden.

Shukla, N., Wickramasuriya, R., Miller, A., & Perez, P. (2015). An approach to plan and evaluate the location of radiotherapy services and its application in the New South Wales, Australia. *Computer methods and programs in biomedicine*, 122(2), 245-256.

Spiegel, P. B., Garber, K., Kushner, A., & Wise, P. (2018). The Mosul trauma response: a case study. *Johns Hopkins Center for Humanitarian Health*, 36.

Toole, M. J., & Waldman, R. J. (1997). The public health aspects of complex emergencies and refugee situations. *Annual review of public health*, 18(1), 283-312.

Toutouh, J., Rossit, D., & Nesmachnow, S. (2019). Soft computing methods for multiobjective location of garbage accumulation points in smart cities. *Annals of Mathematics and Artificial Intelligence*, 1-27.

UN-Habitat. (2016). City profile of Mosul, Iraq. Multi-sector assessment of a city under siege. Accessed 16/9/2019 via [https://reliefweb.int/sites/reliefweb.int/files/resources/UN-Habitat\\_MosulCityProfile\\_V5.pdf](https://reliefweb.int/sites/reliefweb.int/files/resources/UN-Habitat_MosulCityProfile_V5.pdf)

Vafaei, N., & Oztaysi, B. (2014, June). Selecting the field hospital place for disasters: A case study in Istanbul. In *Joint international conference of the INFORMS GDN selection and EURO working group on DSS* (pp. 323-336).

Vahidnia, M. H., Alesheikh, A. A., & Alimohammadi, A. (2009). Hospital site selection using fuzzy AHP and its derivatives. *Journal of environmental management*, 90(10), 3048-3056.

Van der Caaij, T. (2019). Interactively planning resilient and connected field hospital locations in a conflict area. Unpublished MSc thesis. GIMA. University of Twente.

Vasconcelos, J. A., Ramirez, J. A., Takahashi, R. H. C., & Saldanha, R. R. (2001). Improvements in genetic algorithms. *IEEE Transactions on magnetics*, 37(5), 3414-3417.

What We Do. (n.d.) In [Hotosm.org/what-we-do](https://www.hotosm.org/what-we-do). Accessed 18/9/2019 via <https://www.hotosm.org/what-we-do>

WHO [World Health Organisation]. (2008). Integrated Health Services - What and Why? *Technical Brief No. 1*, May 2008.

WMA [World Medical Association]. (2017). WMA statement on medical ethics in the event of disasters. Accessed 16/9/2019 via <https://www.wma.net/policies-post/wma-statement-on-medical-ethics-in-the-event-of-disasters/>

Wu, C. R., Lin, C. T., & Chen, H. C. (2007). Optimal selection of location for Taiwanese hospitals to ensure a competitive advantage by using the analytic hierarchy process and sensitivity analysis. *Building and environment*, 42(3), 1431-1444.

Zeng, L., Krallmann, T., Fiege, A., Stess, M., Graen, T., & Nolting, M. (2019). Optimization of future charging infrastructure for commercial electric vehicles using a multi-objective genetic algorithm and real travel data. *Evolving Systems*, 1-14.