

A decision support framework for large scale emergency response

D. T. Wilson, G. I. Hawe, G. Coates & R. S. Crouch

*School of Engineering and Computing Sciences,
University of Durham, UK*

Abstract

Coordination is often cited as one of the principal problems in large scale emergency response. By representing the problem mathematically, a decision support framework capable of viewing the ‘big picture’ in emergency response, and processing it effectively and efficiently, has been developed. This will result in a provision of high quality suggestions for coordinated response operations, which can be called upon during the decision-making process. In order to ensure the framework is relevant and applicable, significant modeling challenges must be overcome to translate a highly uncertain environment incorporating a complex network of interdependencies into a rigorous mathematical formulation.

Keywords: decision support, mathematical modelling, optimisation, terrorism.

1 Introduction

Following recent high profile and unprecedented large scale emergencies such as the terrorist attacks of September 11th 2001 in USA, or of July 7th 2005 in London, questions relating to how an emergency response operation to such an event should be carried out have gained considerable exposure. Response operations at present are predominantly defined through pre-planned standard operating procedures which, while working well when applied to familiar or foreseeable situations, may not provide the required level of flexibility to ensure a high quality response to unprecedented events.

The provision of a decision support framework, capable of constructing a mathematical model which represents a real-time response operation and applying optimization techniques to find potential solutions, could contribute to ensuring a flexible and robust response strategy. This paper discusses progress made on



the development of such a decision support framework. In Section 2, a selection of key examples of similar work is discussed. Following this, in Section 3 we present the proposed decision support framework at a high level. In Section 4, a case study example will be defined. Details will then be given on the building of a mathematical model of this example in Section 5, before key results are presented in Section 6. Finally, conclusions and further work are discussed in Section 7.

2 Related work

We proceed to briefly review decision support frameworks developed to address response to large scale emergencies, particularly those which feature an optimization model at their center. Firstly, we shall review research which approaches the large scale emergency response problem from the angle of logistics, including those which explicitly incorporate the transportation of casualties in their models. Secondly, we will describe work which addresses the problem in question through a model of resource allocation and task scheduling.

2.1 Logistical formulations

A popular view of the problems presented by large scale emergency response involves consideration of the logistical issues. In particular, it is proposed that response operations could be improved significantly by focussing on the distribution of consumable resources, such as medical supplies, around the emergency environment. These approaches can then build upon the extensive amount of previous research into general vehicle routing and network flow problems.

A decision support model presented in [1] tackles the problem of transporting several different commodities over a network via several possible modes of transport from supply points to demand points. The result of the model is a detailed description of how the available vehicles should move over the transportation network over time along with specifications of the flow of each type of commodity. An approach to a similar problem is presented in [2]. Attempting to incorporate the uncertainty present in real response problems, the model is designed to be used immediately after an earthquake has struck, when information on the epicenter and magnitude has been received but details of actual impact are still uncertain. Decision support is delivered in stages, with the model updating the proposed solution as more information becomes available.

In [3], a model is presented for the inner city transportation of commodities using several modes of transportation following a disaster. The model allows vehicles to be called on from every point in the network, and so can replicate the use of civilian vehicles. Furthermore, it has been designed to facilitate re-planning so as to be of use in the rapidly changing environment of disaster response.

During any urban disaster there will be large numbers of injured civilians that need to be relocated. There has been some effort to incorporate both the transportation of commodities and the collection of casualties into the same model.



In [4] it is noted that helicopters can be used in disaster response operations to both deliver supplies and to transfer the wounded to hospitals, often combining the two tasks into one trip. The model presented is divided into two sub-models, separating tactical and operational decisions. The resulting sub-models present two separate, conflicting objectives (maximizing cost-efficiency for tactical decisions and minimizing total time for operational), with the authors presenting an iterative method for linking the two together.

2.2 Allocation and scheduling formulations

We now examine research which views the problems of interest as some form of *allocation* or *scheduling* problem. [5] provides an example of such work. Here, the authors consider a problem involving the allocation of personnel and equipment to operational areas after an earthquake has struck. The authors identify different types of operational area depending on the type of work that is to be carried out in them, for example search and rescue or stabilization work. The model also accepts as input the survival rates for trapped victims and for casualties who have been rescued but are awaiting treatment. The output of the model consists of a work schedule for all resources which will minimize the expected number of fatalities.

The problem of ambulance dispatch and reallocation during large scale emergency response is tackled in [6]. The authors present the notion of casualty ‘clusters’, suggesting that following a disaster there will be areas with high concentrations of casualties and by dispatching ambulances to these clusters as opposed to individual casualties located far away from any other, a higher level of efficiency will be achieved since ambulances will be able to fill their capacity quickly in each trip. By considering only casualty clusters, the allocation model is significantly reduced in size and is therefore easier to solve in a timely manner.

In [7] a problem of allocating personnel is also addressed, although in this case the units are engineering battalions and the tasks consist of repair work to be carried out following an earthquake. Here the authors decompose the overall problem into two stages, namely assigning a set of tasks to each unit and then describing how each unit should employ its resources to complete the work it has been assigned. The authors note that ‘excessive logistical details’ are omitted since the model is designed for use on a national scale and therefore details like transportation networks are of limited benefit.

2.3 Summary

We have highlighted some common approaches in designing decision support models for use in disaster relief. Various features of the general problem have been noted, from the stochastic nature of demand, through the need for emergency vehicles to re-fuel, to the flexible use of vehicles found in real disaster response situations. The encouraging results presented in the literature suggest that the application of mathematical modeling and optimization to the problems presented in large scale emergency response is worthy of further exploration. However, in



such work it is common for assumptions to be made, some of which may not be appropriate. In [8], the author notes several assumptions commonly made by both practitioners and researchers in the disaster response domain and argues they are invalid. The author concludes that further research in the area could contribute through a more comprehensive modeling approach as opposed to the development of new solution algorithms. Through close collaboration with emergency planning practitioners we can take the opportunity to go some way in this direction. A comprehensive decision support framework with minimal assumptions and more detail can be developed, ensuring those details incorporated are the most pertinent to the problem domain.

3 Decision support framework

A high level overview of the proposed decision support framework is illustrated in Figure 1. It shows the general decision-making process, starting with the gathering of information prior to the event. Following initial information regarding the event, a rule-base may be used to formulate an initial response designed to meet established objectives. An iterative process is then undertaken, comprising of setting objectives, building or updating a mathematical model, and applying solution algorithms to solve it.

4 Problem definition

What follows is a description of a particular large scale emergency, and the characteristics of an associated response, currently being used as a case study for this research. Based on the terrorist attacks in London on July the 7th 2005, the problem describes a scenario involving multiple bombings in crowded areas within an urban environment. We describe in turn three aspects of the problem, namely the environment in which the event takes place, the resources used in the response, and the tasks which are assigned to these resources.

4.1 Environment

A map of London highlighting five disaster scenes, seven hospitals and six fire stations is displayed in Figure 2.

For the purpose of this case study, hospitals may be characterized by two factors, their location and their capacity (in terms of number of available beds). The hospitals shown on the map were used in the actual response to the London bombings - in particular, they were used by the London Ambulance Service – and it is these which will be considered in our mathematical model. The number of casualties taken to each hospital ranged from 6 (Chelsea & Westminster) to 208 (Royal London). In relation to the distribution of casualties during the actual event, the London Assembly report [9] states that “Dispersal of patients to hospitals was uneven because of breakdown in communications within the Ambulance Service.



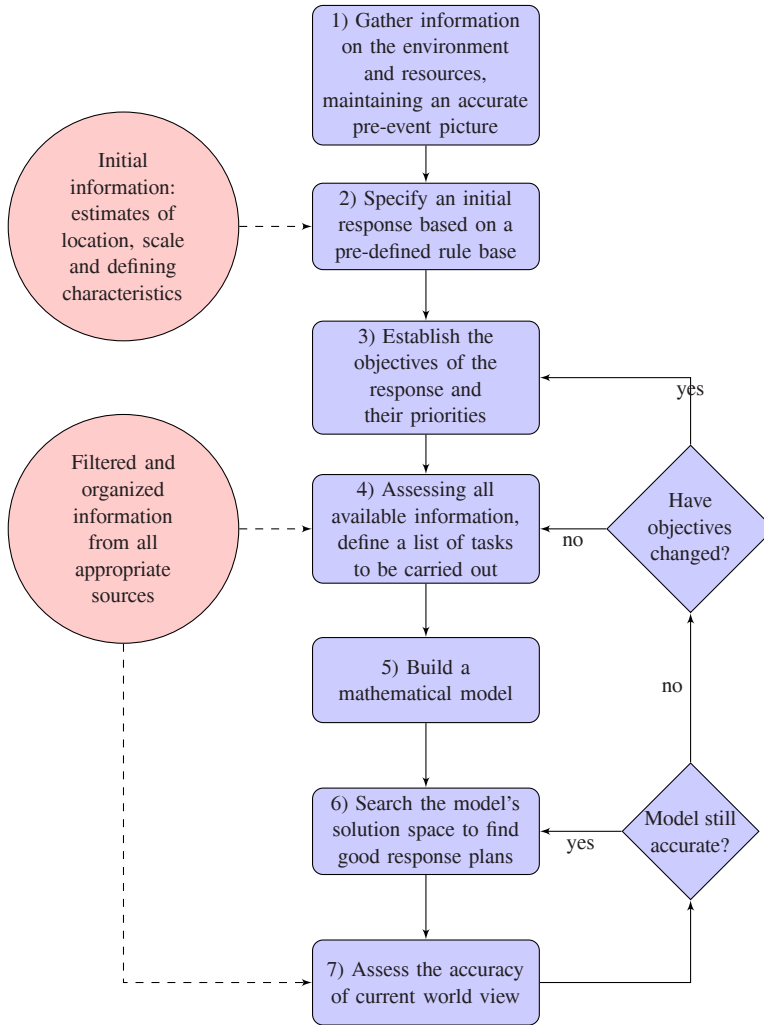


Figure 1: The decision support framework.

In the event, this had minimal impact on patients on 7 July” As such, we may set the capacities of hospitals to be relatively generous in relation to the number of casualties. Specifically, each hospital capacity will equal twenty percent of the total number of casualties under consideration.

The road network of central London may be modeled mathematically as a graph. Nodes of the road network graph represent junctions and points of interest such as hospitals, disaster scenes and fire stations. The edges joining these nodes represent roads, and as such may be described by their distance. Given this graph and the information associated with it, it is then possible to describe routes to be taken by emergency responders and to compute estimates of the resulting travel time.

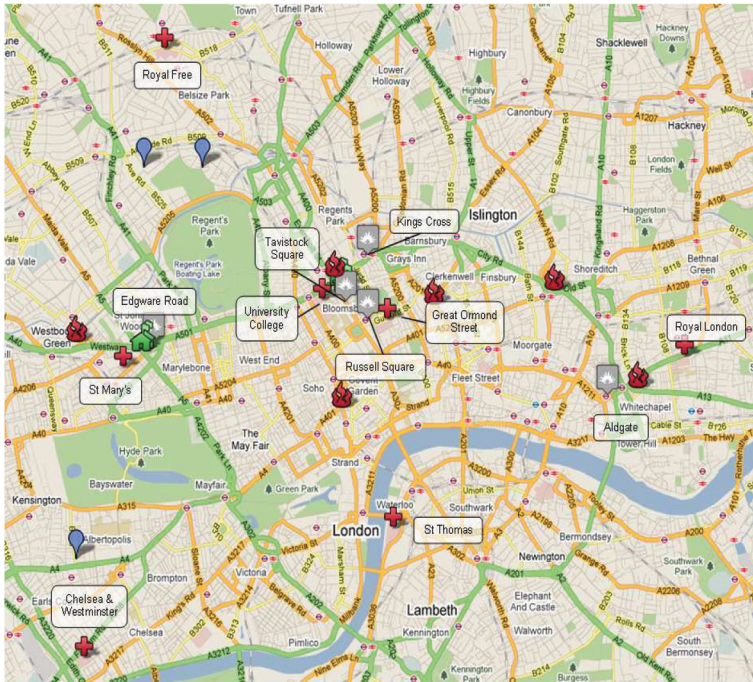


Figure 2: A map of the July 7th bombings. Disaster scenes and hospitals are labeled, while fire stations are denoted by a fire symbol.

Each of the five disaster scenes (Edgeware Road, Tavistock Square, Kings Cross, Russell Square and Aldgate) are described by three parameters: the number of casualties; the severity of injuries; and the stability of the scene. The severity parameter is used to define the distribution of initial health states amongst the casualties at the scene. The parameter will take a value between zero and one, with one corresponding to the most severe scene. The stability parameter is used to define the probability of each casualty being in some way trapped, and so requiring a fire-fighter unit to free them.

Casualties may be described by three parameters - their location, their health and their mobility (i.e. whether they are trapped in some way, and if so to what extent). The health of a casualty is denoted by a discrete classification, namely:

1. treatment may be delayed,
2. treatment is required urgently,
3. treatment is required immediately, and
4. casualty is either dead or death is imminent.

These are the four classifications assigned to casualties during triage [10]. Following the assignment of a health state to a casualty, we expect this health state to change as time passes. How their health changes can depend on what kind

of environment the casualty is in (for example trapped under rubble, in a casualty clearing station or in an ambulance) as well as interventions by medical personnel. Predicting how the health of a casualty is likely to change over the course of a response operation is key to the effective evaluation of proposed response plans.

4.2 Resources

We may think of resources as being divided into four categories - personnel, vehicles, equipment and consumables. In the case of the July 7th bombings, four main groups provided resources to assist with the response operation, namely the health service, the police, the fire and rescue service, and volunteers (both organized and spontaneous). For the purpose of the problem being modeled, we shall focus our attention on the health service and the fire and rescue service.

A key component to the health service response is provided by the Ambulance Service, who provided several types of resources to the response operation of the London bombings. Paramedics represent the backbone of any health service response to a major incident, able to drive ambulances and to administer care to wounded civilians. The training of paramedics and their resulting capabilities can vary - in particular, some paramedics may have received Hazardous Response Team (HART) training in how to operate in an urban search and rescue environment. This particular distinction in abilities is included in our model.

In order to transport casualties from a disaster scene to a hospital, vehicles are required. A two-person paramedic team will have access to an ambulance, capable of carrying one casualty at any one time. For the purpose of this model, we assume that the link between a paramedic and their mode of transport is constant throughout the response operation.

In the scenario under consideration we do not consider the possibility of fires, and so the principle role of the Fire and Rescue service is in the extraction of casualties from within the inner cordon. We consider fire-fighters as being grouped into teams, and that these groupings remain constant over the course of operations. Each fire-fighter team has access to a fire engine vehicle for transportation to disaster scenes, which also carries the equipment necessary for rescue operations.

4.3 Tasks

Large scale emergency response operations may be broken down into a set of individual tasks which can be assigned to resources. In some cases, tasks may need further specification in terms of *how* the resource should go about completing it. These decisions can often be left to the individual resource agents. However, where these decisions may impact the overall response operation, it may be appropriate to make them at the higher level to ensure coordination. Tasks may be divided into two categories, essential and optional. Essential tasks must appear somewhere in some resources schedule in order for the solution to be considered valid.

In the case study under consideration, we consider it essential that each casualty be taken to a hospital. In order for this to happen, each casualty must be taken from



the disaster scene to a designated casualty clearing station. If a casualty is trapped under debris of some kind a *rescue* task is required in order to achieve this. This task may only be performed by a team of fire fighters, and has no pre-requisite for completion. The duration of this task depends on the level of mobility of the casualty in question.

Once a casualty has been rescued, they will join all other casualties at the designated casualty clearing station. From here, each casualty must be transported to an appropriate hospital by a team of paramedics using an ambulance. The duration of this task is dependant on the choice of hospital, given which we calculate the distance of the shortest path and from this estimate the travel time. The decision of which hospital each casualty is to be sent to is necessary at this stage to ensure an effective distribution.

In addition to transporting casualties to hospitals, paramedics may deliver treatment on-scene. This may not always be an optimal course of action and so it is not required to be carried out for every casualty. The duration of the treatment will vary depending on the specific details of the injuries in question. Treatment may be carried out by any paramedic when the casualty is in the casualty clearing station. In addition, paramedics who have received HART training in operating in an urban search and rescue environment may administer treatment to casualties trapped in the inner cordon area who are awaiting rescue by the Fire Service.

5 Mathematical model

An initial mathematical representation of the case study problem is presented in this section. We propose to modify and build upon this initial model over the course of the research project, incorporating more features with each iteration. Combined with regular input and feedback from practitioners, this approach will ensure the end result is as applicable and relevant as possible.

5.1 The solution space

The concept of a solution space is fundamental to the design of a mathematical model. In the context of large scale emergency response, we require a solution to specify how the available resources should be distributed over time and space. Considering the tasks discussed in Section 4.3, we note that one type of task requires further specification when defining a solution – for each hospital trip, we may change which hospital the casualty must be taken to. Given a set of these tasks to be completed, we define a solution to be a set of *schedules*, one for each resource r_i available, of completely specified tasks. Thus, a solution may be partially represented in the form shown in Figure 3.

Combining the above schedules with full specifications for each task defines a solution to the mathematical model. When evaluating a solution of the proposed form the times each task starts and finishes must be known in order to assess the expected outcome of a given solution. We can avoid the incorporation of these times into the solution itself, however, by making the assumption that all resources

Resource	Tasks
r_1	t_1, t_6, t_{13}
r_2	t_3, t_{10}, t_5, t_7
\vdots	\vdots
r_n	t_{11}, t_2

Figure 3: A solution schedule.

will begin work on their next task as soon as possible when taking into account any dependency which exists between tasks. That is to say, given a set of schedules in the form shown above, the start and finish times of each task can be deduced directly. Furthermore, this approach accounts for any variation in the duration of tasks.

5.2 Solution algorithms

Having defined a solution space, we require a method to search for solutions of high quality. For the purpose of this case study, a random local search algorithm was implemented. Given a specific solution \hat{s} , this algorithm proceeds to examine each solution in the *neighborhood* (see Section 5.3 below) in a random order. When a solution s with an improved objective value is found, the algorithm moves to this new solution. The process is repeated until an appropriate stopping criteria is reached.

5.3 The neighborhood structure

Given the above definition of a solution, we must now define a neighborhood structure. This structure helps define how a search algorithm explores the solution space, and is therefore highly influential on any search algorithm's performance. In our initial model we allow two types of neighborhood operations:

1. Moving a task around the schedules,
2. Changing the hospital a casualty is to be taken to.

For example, the algorithm could move from solution s_1 to solution s_2 as illustrated in Figure 4. Since solution s_1 can be converted into s_2 in only one operation, s_1 and s_2 are defined here as neighbors.

5.4 The objective function

Now that we have defined a solution space and specified how a local search algorithm can explore it, we must consider how to evaluate each solution found. Firstly, the start and end times of each task are extracted from the solution. Given these times, the objective function computes a measure of the number of expected fatalities. We wish to model the effect different parts of the disaster environment

	s_1		s_2
r_1 :	t_4, t_1		r_1 : $\mathbf{t_2}, t_4, t_1$
r_2 :	$t_5, t_6, \mathbf{t_2}, t_7$	\rightarrow	r_2 : t_5, t_6, t_7
r_3 :	t_8, t_3		r_3 : t_8, t_3

Figure 4: An example neighborhood operation.

have on a casualties health state. Our problem involves three possible situations a casualty can be in: trapped, waiting to be rescued; free, waiting to be taken to a hospital; and riding in an ambulance. In order to capture these details in our model, a series of discrete time Markov chains was used to estimate how the health of a casualty will change as time passes. We define three Markov chains - each corresponding to a different situation a casualty can be in over the course of the rescue operation. The dynamics of each chain is shown in Figure 5.




Figure 5: Three Markov chains showing injury level changes, where solid lines correspond to the trapped environment, dotted to awaiting transportation, and dashed to being in an ambulance.

Two health states are linked if it is possible to move from one to the other in any given time step. We see, for example, that casualties riding in an ambulance in a injury state of ‘delayed’ may move to state ‘minor’ as they receive basic treatment from paramedics. For the purpose of this case study the actual probability of this transition, and all others, has been estimated.

6 Results and evaluation

To complete the definition of the problem to be considered, we require details of the number of disaster scenes, casualties and resources. We choose these details to mimic the actual emergency of the July 7th bombings. The five disaster scenes and seven hospitals we consider in this problem instance are shown on the map in Figure 2. The number of casualties present at each scene is shown below, together with the expected proportion of casualties in each initial health state (delayed / urgent / immediate / deceased).

As stated in Section 4.1, the capacity of each hospital was set to be twenty percent of the total number of casualties, i.e. 83. The resources available will be spread evenly across the appropriate hospitals and fire stations shown in Figure 2 and will consist of 80 Ambulances, 20 of which are HART trained, and 30 Fire engines. An average speed of 40 miles per hour was used when calculating travel

Table 1: Disaster scene parameters.

Scene	Casualties	Health Distribution	Severity	% Trapped
Edgware Road	58	0.55/0.18/0.14/0.12	0.12	25
Tavistock Square	29	0.15/0.18/0.21/0.47	0.75	1
Kings Cross	143	0.5/0.20/0.16/0.14	0.20	20
Russell Square	54	0.5/0.20/0.16/0.14	0.20	20
Aldgate	128	0.56/0.18/0.14/0.12	0.1	30

times. These figures are designed to approximate those witnessed during the actual London bombings.

Before applying the search algorithm, an initial solution was required. This was created in stages. Firstly, each casualty was assigned to the hospital closest to their initial location, providing that hospital had the required capacity. Secondly, ambulances were assigned a number of casualties they were to take to hospital, and in what order they should be carried out. This was done in such a way as to ensure an even distribution of tasks among ambulances. Furthermore, all casualties with an initial health state of ‘deceased’ were put to the end of their schedule. The allocation of rescue tasks to fire-fighters was carried out in a similar manner. Finally, all optional tasks were omitted from the schedules of this initial solution.

The objective value of the initial solution was calculated to be 116.1. On application of the search algorithm, a solution of objective value 77.1 was obtained, a decrease in the expected number of fatalities of 39. These results suggest there is significant potential for the implementation of the decision support framework to result in an improvement in response operations. Furthermore, qualitative assessment of the framework has been carried out through a number of discussions with emergency planning practitioners from Cleveland Emergency Planning Unit (EPU), Tyne & Wear EPU, Co. Durham & Darlington Civil Contingencies Unit and Government Office for the North East. The feedback received through these discussions has been considered throughout the development of the decision support framework in order to ensure its validity and relevance.

7 Conclusion and further work

By modeling the case study of the London July 7th bombings mathematically, we have been able to capture several pertinent details of the problem. In particular, we have addressed issues arising from the scheduling of a large number of interdependent tasks to a variety of resource types, each with unique capabilities. In addition, we have provided a means to estimate how the health of individual casualties will progress over the course of a response operation, which in turns



allows evaluation of a proposed response operation to be undertaken. Given this mathematical model, we have illustrated the potential of solution algorithms in locating high quality solutions.

Further work will involve the development of the model through the addition of more factors relevant to this case study, before further broadening the scope to include other types of no-notice large scale disasters. In addition, we will develop and implement sophisticated algorithms to ensure high quality solutions can be found in an efficient manner.

Acknowledgements

The authors gratefully acknowledge the funding provided under the UK's EPSRC grant EP/G057516/1. Further, the authors thank practitioners from Cleveland EPU, Tyne & Wear EPU, Co. Durham & Darlington CCU and Government Office for the North East. We would also like to thank Co. Durham & Darlington Fire and Rescue Service, North East Ambulance Service, Northumbria Police, and Tyne & Wear Fire and Rescue Service.

References

- [1] Haghani, A. & Oh, C., Formulation and solution of a multi-commodity, multi-modal network flow model for disaster relief operations. *Transportation Research, Part A*, **30** (3), pp. 231-250, 1996.
- [2] Barbarosoglu, G. & Arda, Y., A two-stage stochastic programming framework for transportation planning in disaster response. *Journal of the Operational Research Society*, **55** (1), pp. 43-53, 2004.
- [3] Ozdamar, L., Ekinci, E. & Kucukyazici, B., Emergency logistics planning in natural disasters. *Annals of Operations Research*, **129**, pp. 217-245, 2004.
- [4] Barbarosoglu, G., Ozdamar, L. & Cevik, A., An interactive approach for hierarchical analysis of helicopter logistics in disaster relief operations. *European Journal of Operational Research*, **140**, pp. 118-133, 2002.
- [5] Fiedrich, F., Gehbauer, F. & Rickers, U., Optimized resource allocation for emergency response after earthquake disasters. *Safety Science*, **35**, pp. 41-57, 2000.
- [6] Gong, Q. & Batta, R., Allocation and reallocation of ambulances to casualty clusters in a disaster relief operation. *IIE Transactions*, **39** (1), pp. 27-39, 2007.
- [7] Brown, G. & Vassiliou, A., Optimizing disaster relief: Real-time operational and tactical decision support. *Naval Research Logistics*, **40**, pp. 1-23, 1993.
- [8] der Heide, E.A., The importance of evidence-based disaster planning. *Annals of Emergency Medicine*, **47** (1), pp. 34-49, 2006.
- [9] The London Assembly report of the 7 july review committee, June 2006. legacy.london.gov.uk/assembly/reports/7july/report.pdf, 2006.
- [10] London Ambulance Service major incident plan. www.londonambulance.nhs.uk/.

