Evaluation of centralised and autonomous routing strategies in major incident response

Duncan T. Wilson *, Glenn I. Hawe, Graham Coates, Roger S. Crouch

School of Engineering and Computing Sciences, Durham University, South Road, Durham DH1 3LE, UK

ABSTRACT

Fast and efficient routing of emergency responders during the response to mass casualty incidents is a critical element of success. However, the predictability of the associated travel times can also have a significant effect on performance during the response operation. This is particularly the case when a decision support model is employed to assist in the allocation of resources and scheduling of operations, as such models typically rely on an ability to make accurate forecasts when evaluating candidate solutions. In this paper we explore how both routing efficiency and uncertainty in travel time prediction are affected by the routing strategy employed. A simulation study is presented, with results indicating that a routing strategy which allows responders to select routes autonomously, as opposed to being instructed via a central decision support program, leads to improvement in overall performance despite the associated increase in uncertainty in travel time prediction.

© 2014 Elsevier Ltd. Published by Elsevier Ltd. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/3.0/).

1. Introduction

A recent study (Zhou et al., 2011) has identified the "application of modern logistics technology" as a critical success factor in emergency management. One element of this application can be seen in the routing of emergency responders during a Mass Casualty Incident (MCI), which has a clear potential to impact on the quality of the overall response operation. This is particularly the case in the response operations of the Ambulance Service, which will involve making many journeys from the affected area(s) to appropriate hospitals. Effective routing decisions in making these journeys will lead to shorter travel times, which will in turn lead to a lower level of suffering and, potentially, a reduction in the number of fatalities.

The use of GPS technology to assist in making effective routing decisions is now commonplace, in the emergency services and more generally. However, the utility of a GPS system may be significantly affected in the period immediately following an MCI, when high levels of disruption (caused directly or indirectly by the incident) can lead to significant uncertainty in the time it will take to travel along a certain route (Jiang et al., 2012). The implication is that a purportedly optimal route specified by a GPS system, based on knowledge of the transport network under regular conditions, may in fact be sub-optimal in the disrupted disaster environment. In such cases it could be argued that routing decisions should be

made with little regard to the guidance offered by the GPS system, with responders instead making decisions themselves based on their prior knowledge of the area and the knowledge acquired as they explore the now disrupted network. This is indeed what happened during the response to the Haiti earthquake, where it has been noted that drivers had "no maps with updated information and had to discover the best routes by driving and exploring" (de la Torre et al., 2012).

Considering the broader problem of resource allocation in MCI response, it has been noted (Altaya and Green, 2006; Simpson and Hancock, 2009) that mathematical models and optimisation algorithms could potentially provide decision support to emergency response personnel, leading to more efficient response operations. However, it is common for such models to rely on an ability to predict the outcome of any given response operation plan. Given this ability, the model can consider a larger decision problem than is feasible for emergency response personnel, accounting for decisions both immediate and in the near future, which in turn allows for better plans to be formulated.

Unknown levels of disruption in a transport network will present a significant challenge to any optimisation model of this type, as it will lead to difficulties in predicting travel times and, consequently, the outcome of the response plan. In this context, it may be beneficial to rely on a centralised specification of routes, acknowledging that the routes themselves may be sub-optimal, in order to improve prediction abilities. If the optimisation model were to release control over routing decisions to the emergency

* Corresponding author. Tel.: +44 113 343 1470; fax: +44 113 343 1471.
E-mail addresses: d.t.wilson@dur.ac.uk, d.t.wilson@leeds.ac.uk (D.T. Wilson).

responders themselves, this would introduce uncertainty over route choice, thereby making the task of prediction more challenging. We hypothesise that this increase in uncertainty, and corresponding reduction in utility of the decision support program, could eclipse any benefit gained through better routing and subsequently shorter travel times.

In this paper we build on previous work which has introduced a scheduling-based decision support program for MCI response and demonstrated its sensitivity to uncertainty in temporal parameters (Wilson et al., 2012). Having previously ascertained that disruption to the transport network can be a significant source of such uncertainty (Wilson et al., 2013), in this paper we describe and evaluate routing policies designed to mitigate against these problems.

1.1. Routing in decision support systems for MCI response

Transport networks are not always explicitly modelled within decision support programs designed for disaster response. For example, Wex et al. (2012) present a scheduling model designed to assist in the allocation of response units to incidents, taking as input the travel times associated with each possible journey response units may make. Similarly, a travel time matrix describing the relation between points of interest is taken as problem input by Zhang et al. (2012) in the model of emergency responder allocation. This can be contrasted with work such as that of Yi and Kumar (2007) and Haghani and Oh (1996), where the transport network is represented as a graph, with each edge assigned a parameter describing the time needed to traverse it. Where such graphical representations of transport networks are included, it is common to assume their structure and parameters are deterministic and constant over time. This is true both of models designed to assist in commodity distribution over a large geographic area, such as those presented in Chang et al. (2007), Sheu (2007), Lin et al. (2011), and Tzeng et al. (2007), decision support programs using a scheduling formulation to assign tasks to emergency responders (Rolland et al., 2010; Wilson et al., 2012), and routing based formulation for the support of casualty transportation and evacuation (Chiu and Zheng, 2007; Yi and Ozdamar, 2007). By assuming all necessary information regarding the transport network is readily available, routing decisions can be made with confidence using a standard shortest path algorithm.

It is common in past work to use a reduced simplification of the actual transport network when representing it as a graph, an approach which can help avoid excessive computational burden. In the problem scenarios considered by Yi and Kumar (2007), for example, the most complex network considered contains 80 nodes connected by 1600 edges. Considering a geographic area large enough to encompass six cities in Turkey, the model presented by Ozdamar et al. (2004) represents the transport network using 12 nodes and 12 links, based upon motorway infrastructure. In contrast, a dense network comprised of 34,890 nodes and 43,445 links is used in the test problem considered by Jotshi et al. (2009), with a hierarchial decomposition employed to assist route computation in a timely manner.

Uncertainty in the disruption of the transport network has been incorporated to a limited extent using stochastic programming formulations. Examples include (Barbarosoglu and Arda, 2004; Mete and Zabinsky, 2010; Rawls and Turnquist, 2010), which consider a finite number of scenarios, each with assigned probability and associated network parametrization. Uncertainty is also acknowledged in the work of Jotshi et al. (2009), which extends the ambulance allocation model presented by Gong and Batta (2007) by including a data fusion step to estimate the level of damage and disruption on each road link. A solution methodology for finding optimal paths in a disrupted network following a disaster is presented in Zhang et al. (2013). The authors employ the network

representation described by Yuan and Wang (2009), where the travel time associated with each edge of the transport network is assumed to increase over time in a manner which reflects its proximity to the disaster. A dynamic transport network structure is also modelled in the work of Fiedrich et al., 2000, with nodes and edges being added or taken away to reflect the impact of both the disaster and the response operation.

1.2. Contribution of this paper

In recent reviews of optimisation models for emergency logistics (Caunhye et al., 2012; de la Torre et al., 2012) it has been noted that there has been little research in the area employing stochastic models. Given the potential for an MCI to disrupt the transport network, directly or indirectly, and thus lead to uncertainty in routing and travel time prediction, this is clearly an area which merits further research. While some authors have acknowledged the possibility of disruption to the network and the subsequent uncertainty, it remains unclear whether or not this uncertainty will ultimately reduce the utility of a decision support program, and how any such effect depends on the choice of routing policy.

The remainder of this paper will be structured as follows. In Section 2 we will briefly describe a previously published scheduling based decision support program designed to optimise resource allocation in MCI response. In Section 3 we go on to present a simulation routine designed to generate random levels of disruption to the transport network representative of the problem environment. A number of potential routing policies will be introduced in Section 4, with details provided on the associated simulation of route choice and prediction of travel time. These policies will be compared using a Monte Carlo approach in Section 5, allowing for the uncertainty in the problem to be fully captured.

2. A scheduling model for disaster response

In this paper we employ the multi-objective optimisation model described in Wilson et al. (2013). The model is of a task scheduling nature, similar to the Flexible Job Shop scheduling Problem (FJSP) (Brandimarte, 1993). Specifically, each casualty in the problem is associated with a number of tasks which must be carried out by the available emergency responders. The tasks associated with each casualty will always include a *transportation* task, which requires an *ambulance* responder and involves the transportation of the casualty from the incident site i to a chosen hospital h. Other tasks include *treatment* tasks and *rescue* tasks, and have a specific order in which they must be carried out in. This leads to a dependency structure in the scheduling model. A solution is defined by an ordered allocation of tasks to emergency responder units, together with a mapping from the set of casualties to the set of hospitals. Given such a solution, the first stage in its evaluation is the creation of a corresponding schedule by estimating the time at which each task will start and finish. This involves estimating the duration of each task, respecting the dependency relations which exist between them, and estimating associated travel times.

An example segment of a response schedule is given in Fig. 1, where the initial schedule of two responders r_1 and r_2 are shown. In addition to displaying the tasks to be carried out, movement between different areas in the MCI environment are shown.

As can be seen in Fig. 1, the accurate estimation of travel times is an essential part of computing an accurate schedule. The objective functions which measure the quality of a given schedule primarily use the estimated start and end times of tasks in their computations, implying that the accuracy of travel time estimation will directly affect the model's ability to accurately compare solutions and select one of high quality.
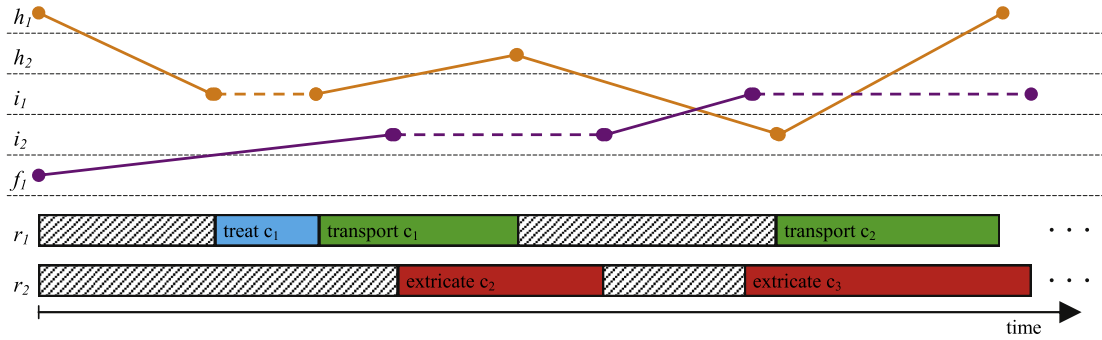
Fig. 1. An extract from a simple schedule involving two responder (r_1, an ambulance, and r_2, a fire appliance) processing two casualties (c_1 and c_2), involving travel between a fire station (f_1), two hospitals (h_1 and h_2) and two incident sites (i_1 and i_2).

2.1. Objective functions

Schedules are evaluated using objective functions denoted f_1 and f_2, where the former predicts the number of fatalities resulting from the specified schedule and the latter the level of suffering. In order to predict the number of fatalities arising during a response operation, the model uses information regarding the initial health level of each casualty (as assessed during a triage operation (Advanced Life Support Group, 2011)) together with the time taken to remove that casualty from the dangerous incident site environment and take them to a safe environment where they can receive treatment and subsequent transportation to hospital.

The measurement of suffering, f_2, is determined by the time taken to transport casualties to hospital, weighted by their health level. Consideration is also give to the suitability of hospital choice, in terms of capacity and capabilities. The two measures are combined in a lexicographic manner, with the full optimisation model defined as

$$\min_{s \in S} f_1(s), \quad f_2(s). \tag{1}$$

The specific form of these functions is omitted in this paper, as the focus is on routing policies and their effect on the utility of the optimisation model. Full details can be found in Wilson et al. (2013).

2.2. Online optimisation

The model described in Wilson et al. (2013) has been extended to allow for use in an *online* manner. This is in contrast with the default *offline* usage. In the offline case the model is initialised upon collecting all necessary information, a local search optimisation algorithm is applied and allowed to run for a short period of time, after which the resulting solution is taken as the response operation plan and the responder units are instructed accordingly.

In contrast, the online approach uses the local search procedure is used to search the solution space in real time. At a point where a responder finishes an allocated task, the best schedule found by that point in time is consulted to find which task the responder should be given next. The search process then continues, noting that this task has been issued and is therefore no longer a component of the solution space. Regarding travel times, information is passed back to the model when a responder sets out on a journey and again when they arrive at their destination. As such, a regular supply of travel time data is received as the response operation progresses. This allows, potentially, for any predictions of future travel times to be revised and improved in real time. This possibility will be explored in terms of each routing strategy considered in Section 4.

2.3. Travel times

Prior to detailing how transport network disruption is modelled it is helpful to describe the general procedure which will be used to estimate the travel time of a responder along a route of specified distance. Given a specific route with total distance d, travel times are estimated using the model described by Kolesar et al. (1975), as recently validated by Budge et al. (2010). The function, denoted $KWH(d)$, gives an estimate of the median travel time for that route. In order to find d, Dijkstra's algorithm (Skiena, 1990) is applied to the transport network graph. The median travel time is then estimated as

$$\hat{m} = KWH(d) = \begin{cases} 2.42\sqrt{d}, & d \leqslant 4.13 \text{ km} \\ 2.46 + 0.596d, & d > 4.13 \text{ km} \end{cases}, \tag{2}$$

where $4.13 = v_c^2/2a$ denotes the distance required to travel in order to reach "cruise speed" v_c and a is the average acceleration of the vehicle as it increases speed to v_c. The values of these parameters are taken from the analysis of ambulance travel times in Calgary, Canada, presented in Budge et al. (2010).

3. Transport network disruption

As discussed in Section 1, in this paper we will consider scenarios where the transport network has been in some way disrupted, either directly or indirectly, by the MCI. In this case we consider network disruption to be represented by longer travel times associated with the individual edges comprising the network. In this section we provide further details regarding how such a disruption may be simulated, and analyse the empirical properties of this simulation model by examining resulting travel times and optimal routes.

3.1. Method

In order to transform the standard transport network, denoted G, into one which has been disrupted, denoted G^*, we modify the distance parameter of each road link. Specifically, a random variable $Y \sim exp(\lambda)$ is sampled for each link, the distance of which is then multiplied by the factor $(1 + Y)$. We can interpret λ as a parameter representing the level of disruption to the transport network. For example, setting $\lambda = 0.5$ will lead to the distance parameter of each link on the road network being increased on average by a factor of $(1 + E[Y]) = (1 + 2) = 3$.

3.2. Empirical properties

The purpose of this simulated disruption is to generate uncertainty regarding travel times and optimal routing within the

transport network, and not to realistically simulate the effects of a MCI on the network. Whilst such a realistic simulation model would be desirable, we will show that the approach described does indeed generate uncertainty in both travel time estimation and optimal route choice and is therefore fit for our purpose. To demonstrate this, we consider a specific journey in the problem scenario described in Section 5.1, namely a journey from a hospital to an incident site. Under normal conditions the median travel time for the journey can be calculated as 2.92 min using Eq. (2), where the shortest path was calculated using Dijkstra's algorithm on the standard network parameterization.

To gain an understanding of the effect of our disruption process, we simulated 500 instances of disruption with $\lambda = 0.5$ and calculated two quantities for each simulation run. Firstly, the shortest path calculated by applying Dijkstra's algorithm to the standard network G, denoted p, was found. The 'distance' of this path *in the disrupted network*, defined as the sum of the distance parameters of each of its edges, is calculated and denoted $D_G(p)$. The associated estimate of the median travel time, denoted $\hat{m}$, is calculated using Eq. (2): $\hat{m} = KWH(D_{G^*}(p))$. This estimated travel time $\hat{m}$ corresponds to a route which was chosen based on prior knowledge of the transport network, with no information regarding the disruption of the network.

Secondly, the actual shortest path over the disrupted network G^*, denoted p^*, was found. Again, the estimated median travel time along this path was calculated, $\hat{m}^* = KWH(D_{G^*}(p^*))$. Here, $\hat{m}^*$ is the estimated travel time for the actual shortest path, as calculated with full knowledge of the disruption to the network. The difference in travel times between the route chosen based on prior knowledge and the route chosen based on perfect knowledge, $\hat{m} - \hat{m}^*$, represents the utility of perfect information regarding the disruption of the network. That is, if one were to have access to information describing how the network has been disrupted, a travel time saving of $\hat{m} - \hat{m}^*$ could be made due to being able to find the actual shortest path on the disrupted network.

The results of this analysis are presented as a scatter plot with marginal histograms in Fig. 2. It is clear that the disruption has a large and variable effect upon the median travel time across route p, $\hat{m}$, as can be seen in the corresponding histogram. The median travel time along the actual shortest path, p^*, is typically lower, with an average difference of 1.19 min. We can therefore conclude that the disruption process described achieves both of our aims,

Fig. 2. The joint and marginal empirical distributions of travel times on a disrupted network using both pre-calculated (p) and actual (p^*) shortest path routes.

namely: increasing the uncertainty in the travel time associated with a specific route, and creating alternative routes with a shorter median travel time than the pre-computed shortest path.

4. Routing policies – simulation and prediction

Having established that the proposed simulation of network disruption does introduce uncertainty in both travel time and optimal route choice, it remains to consider how to modify the optimisation model of Wilson et al. (2013) to adapt to this challenge. In this paper we will describe and evaluate four routing strategies which have the potential to assist in this manner, where each one specifies the simulation and prediction of travel times. Two strategies (Static Routing and Centralised Adaptive Routing) are of a centralised nature, in the sense that the routing choice of each journey is made by the central optimisation model. The remaining two strategies (Autonomous Individual/Collective Adaptive Routing) are of an autonomous nature, where routing decisions are made by individual responders without recourse to the central model.

4.1. Static routing (SR)

For any given journey (as defined by a pair of locations) a single route is specified centrally at the outset of the response operation and is used for the duration by all responders. Thus, this policy employs a network reduction approach, reducing the full transport network graph to a simplified matrix representation connecting each point of interest to be used throughout the response operation. This policy will allow for relatively accurate travel time prediction as there will be no uncertainty in route choice, but will likely involve relatively poor routing decisions. This fact was demonstrated in Fig. 2 of Section 3, where our analysis showed an average difference of 1.19 min between median travel times of paths p and p^*.

Whereas the scheduling model uses the median travel time of a given route in its predictions, in a simulation of a response operation the actual travel times will be subject to random variation around this median. As discussed in the work of Westgate et al. (2011), travel times may be modelled as random variables X following a lognormal distribution,

$$X \sim \log N(\mu, \tau), \tag{3}$$

with an assumed $\tau = 0.00227$ (following (Westgate et al., 2011)). Under this routing policy the route p will be determined by the optimisation model and will therefore be known, allowing the distance to be calculated as $d = D_{G^*}(p)$. Noting that the median of the lognormal distribution is given by $m = e^\mu$, the parameter μ can be calculated using the median value obtained in Eq. (2), $m = KWH(d)$, following which a variate X can be drawn. This process is illustrated in Fig. 3.

Given the real-time, online nature of the problem as described in Section 2.2, we employ a Bayesian approach in revising the estimate of the unknown parameter μ as more travel time data becomes available. Specifically, using the conjugate prior distribution for μ under the lognormal likelihood,

$$\mu \sim N(\mu_0, \tau_0), \tag{4}$$

we can calculate the posterior distribution following the observation of n data x_i,

$$\mu \sim N(\mu_n, \tau_n), \tag{5}$$

where

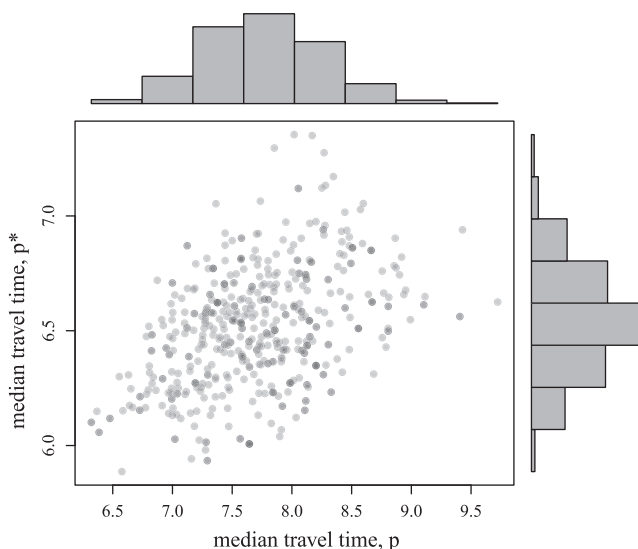$$\mu_n = \frac{\tau_0 \mu_0 + \tau \sum_{i=1}^{n} \ln(x_i)}{\tau_0 + n\tau} \tag{6}$$
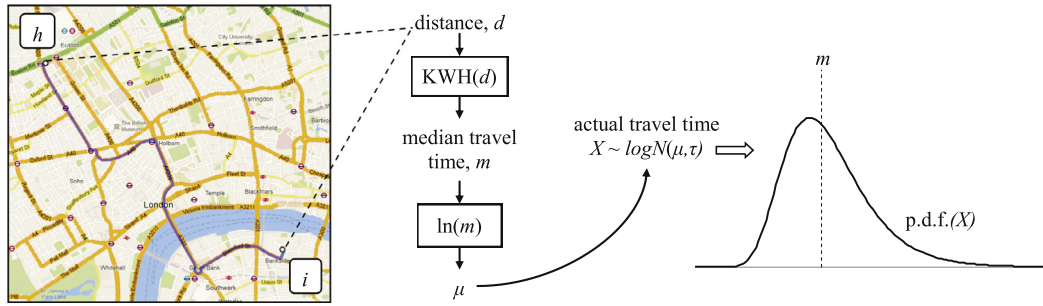
Fig. 3. The process whereby a travel time can be sampled given a route of distance d.

and

$$\tau_n = \tau_0 + n\tau. \tag{7}$$

The expectation of this posterior distribution, $\hat{\mu} = E(\mu)$, is then used as an estimate of μ, giving $X \sim log N(\hat{\mu}, \tau)$. As noted previously, the median travel time for the route in question can then be estimated as $m = e^{\hat{\mu}}$. This routine is carried out for each single travel time data x immediately upon its observation. As such, our ability to predict travel times will improve as the response operation continues and more travel time data are observed.

4.2. Central adaptive routing (CAR)

At the start of each journey a route is calculated centrally using the current perception of the transport network, denoted G_i. This perception is modified each time a journey is completed and the travel time information is communicated to the optimisation model, giving a new network parameterization G_{i+1}. This policy may improve upon the routing decisions of SR as it allows for routes alternative to the initial route p to be explored, at the cost of increased error in travel time prediction.

The process used to update the current perception of the transport network, G_i, upon receiving information regarding a travel time for a particular route p_i, employs a heuristic which adjusts the 'distance' parameters of all links in the route p_i by a factor determined by comparing the realised travel time t with the predicted travel time $\hat{t}$. Specifically, the distance parameter of each link is multiplied by the factor $t/\hat{t}$. This results in the updated network parameterization G_{i+1}, which is used in subsequent routing decisions in conjunction with Dijkstra's shortest path algorithm.

As in the case of SR the routing decisions are known, having been issued by the central DSP. We can therefore use the same routine to generate a sampled travel time using the information contained within the disrupted network representation G^*. In terms of travel time prediction, given the transport network representation G_i used to find the route in question, p_i, the travel time is estimated to be the median time according to Eq. (2), $m = KWH(D_{G_i}(p_i))$.

4.3. Autonomous Adaptive Routing (AIAR and ACAR)

In Autonomous Individual Adaptive Routing (AIAR) each responder makes their own routing decisions upon making a journey. They do so in an isolated manner with no communication with either the central optimisation model or other responders, and learn from past experience such that their routing choices improve on average as more journeys are completed. By introducing uncertainty in route choice, over and above that arising from disruption to transport network parameters and standard travel time variance, this policy will lead to larger errors in travel time prediction than in SR and, possibly, CAR. Autonomous Collective Adaptive Routings (ACAR) is identical to AIAR but with information being

freely shared among responders, allowing the learning process to be done in a collective manner, resulting in faster convergence towards the optimal routing decisions.

In order to simulate the natural improvement in routing choices which would be made by a responder, or set of responders under the ACAR policy, we use a Markov process to generate a sequence of true median travel time values which can then be used, as illustrated in Fig. 3, to define lognormal distributions from which travel times can be sampled.

We assume that the first route chosen by a responder will be the shortest route under normal conditions, p. Thus, we can use the disrupted road network to find the initial median travel time, m_1. The travel time of the first trip is then sampled as $X_1 \sim log N(\ln(m_1), \tau)$. In order to simulate the next median travel time, m_2, we use from a normal distribution with mean αm_1 and variance β^2. This is then used when sampling the actual travel time $X_2 \sim log N(\ln(m_1), \tau)$. The process continues until a median travel time less than or equal to the best possible route, under the disrupted network, has been reached. The process is illustrated in Fig. 4, which shows how the simulated median travel times m_i and the actual travel times x_i progress over ten journeys.

Also shown in Fig. 4 is the estimated travel time for each journey. This is the value used by the optimisation model when predicting all future journeys of that type. Whereas under policies SR and CAR the routing choice was known to the optimisation model, in the case of autonomous routing this is not the case. Accordingly, making accurate predictions of travel times is more challenging. We propose using an exponential smoothing method when predicting the travel time of the ith journey based on past observed travel times. Specifically, denoting the travel time of journey i as t_i,

$$t_{i+1} = \gamma \times t_i + (1 - \gamma) \times t_{i-1}. \tag{8}$$

In the experimental analysis presented in this paper we will use a smoothing factor of $\gamma = 0.5$.

In addition to the smoothing factor, two other parameters are required to specify the nature of the simulation: α and β^2. The average improvement in travel time achieved through better routing at each iteration is given by α, while β^2 controls the amount of random variation around this average. Without a comprehensive source of data relating to travel times in a MCI environment it is not possible to derive empirical estimates for these parameters. In the experiments described in Section 5, the value of β^2 will be held constant throughout, whereas the value of α will be adjusted to examine the effect on overall performance.

5. Experiments and analysis

In order to achieve the aim of this paper, i.e. to evaluate a number of potential routing policies in terms of both their ability to find short routes and their effect on the predictability of temporal
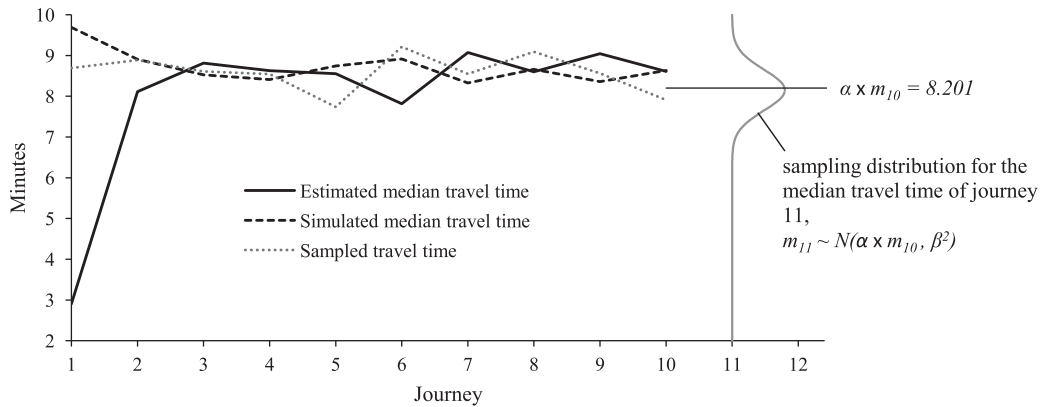
Fig. 4. An instance of route choice progression under the *ASAR* policy.

parameters within the optimisation model, we present a Monte Carlo experimental analysis. In this section an example MCI scenario will be defined, and this will be used for the throughout the analysis. Several variables including the numbers of casualties and responders remaining constant, allowing for effects of variation in the parameters explicitly related to the routing problem to be explored explicitly.

5.1. Case study

We consider a scenario involves one, two or three separate incident sites across central London, UK, with a total of 210 casualties distributed evenly across all sites. The response resources available consist of 53 ambulances (with crew) and 27 fire appliances (with crew). The environment includes a graph representing the central London road network at a fine level of detail (with 21,214 nodes and 29,225 edges). In solving the problem using the optimisation model describe in Section 2, a number of tasks relating to each casualty (namely their extrication, treatment and transportation to hospital) must be assigned to appropriate responders and ordered in such a way as to minimise the objectives f_1 and f_2.

Fig. 5 illustrates the locations of the three (potential) incident sites along with the three nearby hospitals to which casualties may be transported. Over the course of the operation responders will frequently be instructed by the DSP to travel not only from incident site to hospital, but also between incident sites. As such, there are nine (one incident site), ten (two incident sites) or twelve (three incident sites) journeys for which routing choices must be made and travel times estimated according to the number of incident sites. The transport network of the illustrated area is represented by a high-resolution graph based on data provided by Ordnance Survey MasterMap, which was processed using the STORMI package described in Hawe et al. (2012).
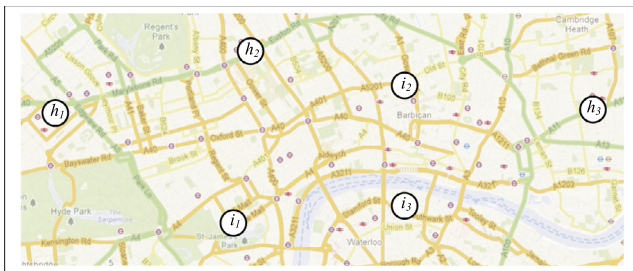

Fig. 5. Three incident sites and three hospitals in central London, as part of the test problem environment.

5.2. Experimental design

A number of factors exist which may affect the performance of the various routing policies. In particular, we are interested in:

1. The parameter describing the disruption to the network, λ (Section 3).
2. The parameters describing the ability of responders to autonomously search for high quality routes, α and β^2 (see Section 4.3).
3. The number of sites which comprise the MCI, S_{one}, S_{two} and S_{three}.

Our principle goal in the analysis that follows is to compare the performance of autonomous routing policies *ASAR* and *ACAR* with a baseline policy *SR* and an alternative *CAR*. This will include determining which values α will lead to autonomous policies outperforming centralised policies. That is, how well must responders be able to route themselves to justify removing the routing decision from the optimisation model? Answering this question will inform the design of future optimisation models for MCI response. The variation of both disruption level and the number of sites is important to explore how this relationship varies with underlying problem characteristics.

Given the several sources of uncertainty within the model, a Monte Carlo approach is employed. That is, for any given point in the experiment space, n instances of the problem are generated and solved in order to estimate the distribution of the final objective values. The points within the experiment space are defined through a standard factorial design based on the following factors:

- Routing policy – {SR, CAR, AIAR, ACAR}.
- Road disruption, λ – {2, 1, 0.5, 0.25}.
- Number of incident sites – {1, 2, 3}.

This gives a total of $4 \times 4 \times 3 = 48$ experimental design points, where each point corresponds to a unique combination of the three factors. For the purposes of the initial evaluation, the autonomous improvement parameter pertaining to the routing policies was set as $\alpha = 0.95$.

Following this initial set of experiments, a second set was designed to focus on the effect of altering the α parameter governing the rate of improvement in the autonomous routing policies. The points within the experiment space are defined through another standard factorial design based on the following factors:

- Routing policy – {AIAR, ACAR}.
- Road disruption – {2, 1, 0.5, 0.25}.
- Number of incident sites, λ – {1, 2, 3}.
- Autonomous improvement, α – {0.9, 0.8}.

This gives a total of $2 \times 4 \times 3 \times 2 = 48$ experimental design points. These results were combined with those of the previous set of experiments, considering only the routing policies AIAR and ACAR. This provided a total of three levels of α in total – 0.95, 0.9 and 0.8.

5.3. Results

The results of the first set of experiments are summarised in Tables 1 and 2, which focus on the objectives of fatalities (f_1) and suffering (g_2) respectively.[1] Each table provides the average objective value obtained across all experiments for a given pair of road network disruption levels (denoted by λ) and routing strategy. In order to indicate the level of variability around these average values, the standard deviation is also provided.

As can be seen in Table 1, the routing strategies SR (static routing) and CAR (central adaptive routing) lead to similar performance in terms of the fatalities objective, f_1. This is consistent across each of the four levels of road network disruption considered. Moreover, with the rate of improvement set to $\alpha = 0.95$, similar performance is also observed when using the routing strategy ACAR (Autonomous Collective Adaptive Routing). As would be expected, the strategy AIAR leads to worse performance in comparison to ACAR. This is due to the fact that in ACAR the improvements in routing are designed to reflect a sharing of information amongst responders, whereas in AIAR each responder works independently to improve their routing choice.

Table 2 presents corresponding summary statistics in relation to the suffering objective, g_2. It is clear that the relation between routing strategies broadly mirrors that of the f_1 case, with routing strategies SR and CAR leading to similar performance, both significantly better than that obtained through strategy AIAR. However, in this case the routing strategy ACAR appears to offer some benefit over strategies SR and CAR. This benefit diminishes as the extent of disruption to the road network decreases.

The results presented in Tables 1 and 2 are broken down by routing strategy and level of network disruption, but averaged over the number of sites comprising the MCI. In order to assess the extent to which this factor influences performance, linear regression models were fitted to the data. Specifically, two models were constructed, one for each of the objectives of interest. Included as predictor variables were the number of sites, the level of disruption and the routing strategy. Each variable was taken to have a qualitative, or categorical, nature. This is the natural representation for the routing strategy, but the quantitative nature of both the number of sites (which can take values 1, 2 or 3) and the level of disruption (which can take values 0.25, 0.5, 1 or 2) suggests that a numerical representation may be the natural choice for these variables. However, such a representation would correspond to assuming a linear relationship between the predictor and the dependant variable. A qualitative representation does not require this assumption, and therefore provides greater flexibility.

The resulting regression models are described in Table 3, where estimates of the effects of each predictor variable are given along with their standard error. The reference values are taken as S_{one}, $\lambda_{0.25}$ and ACAR for the number of sites, road network disruption and routing strategy respectively. Asterisks denote to what extent the effects are judged to be statistically significant (i.e., the extent to which the observed effect is unlikely to be due to chance alone).

As shown by the high value of the adjusted R^2's, both models fit the data well and as such no further terms, such as quadratic or interaction terms, were added. The models confirm what was

suggested in the summaries of the data given in Tables 1 and 2, as they show that the choice of routing strategy has no significant effect upon the fatalities objective (discounting the choice of AIAR), but does have a significant effect upon the suffering objective.

The above analysis used a value of $\alpha = 0.95$ for the rate of improvements in routing choice in the autonomous routing strategies, AIAR and ACAR. As described in Section 4, this means that every time a responder makes a specific journey, the length of the route chosen will be (on average) 0.95 times the length of the last route chosen for that same journey. Using this parameter value, the routing strategy ACAR was shown to lead to comparable performance to the centralised routing strategies, SR and CAR. It remains to be seen to what extent lower values of α will lead to the strategy ACAR out-performing SR and CAR.

In order to investigate this, the second set of experiments described at the beginning of Section 5.2 were run. The resulting data set was combined with the previous data, together allowing for comparisons in the performance of strategy ACAR for three values of α, 0.95, 0.9 and 0.8. The results are summarised in Tables 4 and 5, which show the average (standard deviation) values of

Table 1
Average (standard deviation) fatalities objective values under each routing strategy, for varying levels of disruption.

	$\lambda = 0.25$	$\lambda = 0.5$	$\lambda = 1$	$\lambda = 2$
SR	67.3 (2.9)	56.1 (1.8)	49.4 (1.7)	45.6 (1.8)
CAR	67.8 (2.4)	56.2 (2.1)	49.1 (1.8)	45.6 (1.9)
AIAR	80.1 (2.7)	65.6 (1.9)	55.6 (1.8)	49.7 (1.6)
ACAR	68.1 (2.3)	56.0 (1.6)	49.1 (1.8)	45.5 (1.8)

Table 2
Average (standard deviation) suffering objective values under each routing strategy, for varying levels of disruption.

	$\lambda = 0.25$	$\lambda = 0.5$	$\lambda = 1$	$\lambda = 2$
SR	19,750 (1251)	14,840 (889)	12,048 (630)	10,670 (581)
CAR	19,589 (1393)	14,904 (983)	12,035 (673)	10,608 (521)
AIAR	23,263 (1679)	17,050 (962)	13,305 (646)	11,153 (603)
ACAR	19,164 (1243)	14,517 (803)	11,832 (636)	10,584 (591)

Table 3
Fitted linear regression models for fatalities and suffering objectives, examining the effect of routing strategies.

	Dependent variable	
	f_1	g_2
S_{two}	0.154	795.390***
	(0.128)	(48.273)
S_{three}	1.137***	1457.455***
	(0.128)	(48.280)
$\lambda_{0.5}$	−12.389***	−5141.224***
	(0.148)	(55.524)
λ_1	−20.058***	−8155.442***
	(0.149)	(56.181)
λ_2	−24.261***	−9692.291***
	(0.148)	(55.729)
AIAR	8.109***	2180.700***
	(0.148)	(55.748)
CAR	−0.025	249.071***
	(0.148)	(55.775)
SR	−0.052	321.971***
	(0.149)	(55.939)
Constant	68.399***	19013.970***
	(0.157)	(59.016)
R^2	0.945	0.948
Adjusted R^2	0.944	0.948

Note: *$p < 0.1$.
**$p < 0.05$.
*** $p < 0.01$.

---

[1] To obtain the data presented an analysed in this paper, please contact the lead author.

Table 4
Average (standard deviation) fatalities objective values under autonomous routing strategies, for varying levels of disruption.

	$\lambda = 0.25$	$\lambda = 0.5$	$\lambda = 1$	$\lambda = 2$
$\alpha = 0.95$	68.1 (2.3)	56.0 (1.6)	49.1 (1.8)	45.5 (1.8)
$\alpha = 0.9$	68.1 (2.7)	56.7 (2.1)	49.0 (1.9)	45.5 (1.6)
$\alpha = 0.8$	67.3 (2.1)	56.2 (1.9)	49.2 (1.8)	45.1 (1.8)
SR	67.3 (2.9)	56.1 (1.8)	49.4 (1.7)	45.6 (1.8)

Table 5
Average (standard deviation) suffering objective values under autonomous routing strategies, for varying levels of disruption.

	$\lambda = 0.25$	$\lambda = 0.5$	$\lambda = 1$	$\lambda = 2$
$\alpha = 0.95$	19,163 (1243)	14,517 (803)	11,833 (636)	10,584 (591)
$\alpha = 0.9$	18,949 (1194)	14,396 (686)	11,779 (627)	10,639 (559)
$\alpha = 0.8$	18,818 (1120)	14,383 (809)	11,758 (725)	10,443 (596)
SR	19,750 (1251)	14,840 (889)	12,048 (630)	10,670 (581)

objectives f_1 and g_2, respectively, factored by the level of road network disruption and value of α.

Table 4 shows that, in terms of objective f_1, altering the parameter α in the manner described does not lead to any significant change in the performance of the optimisation model when using the routing strategy ACAR. However, some difference is observed in objective g_2. The benefit afforded through the routing strategy ACAR observed when $\alpha = 0.95$ (see Table 2) is further enhanced as the learning rate is improved. In the most extreme case considered, where $\alpha = 0.8$ and $\lambda = 0.25$, the strategy ACAR leads to 4.7% better performance in terms of g_2 in comparison to strategy SR.

Similarly to the previous analysis conducted on the first set of experiments, linear regression models were also fit to the data, allowing for different values of α in order to further quantify its effect. Considering only the routing strategy ACAR, models were fitted to both fatalities and suffering objective values using the number of sites, the level of disruption and the rate of autonomous improvement, α. As in the models described in Table 3, number of sites and level of disruption were coded as qualitative variables. The rate of autonomous improvement was coded as a quantitative variable. The resulting models are described in Table 6.

Again, the models are judged to fit the data well, as demonstrated by the high values of adjusted R^2's, and so no further terms

Table 6
Fitted linear regression models for fatalities and suffering objectives for model, examining autonomous improvement rate.

	Dependent variable	
	f_1	g_2
S_{two}	−0.058	763.799***
	(0.167)	(55.259)
S_{three}	1.042***	1,403.707***
	(0.167)	(55.044)
$\lambda_{0.5}$	−11.832***	−4,613.958***
	(0.192)	(63.389)
λ_1	−18.827***	−7,262.311***
	(0.193)	(63.815)
λ_2	−22.507***	−8,497.965***
	(0.192)	(63.471)
α	1.163	1,169.000***
	(1.256)	(414.328)
Constant	66.548***	17,279.780***
	(1.166)	(384.609)
R^2	0.955	0.967
Adjusted R^2	0.955	0.966

Note: *$p < 0.1$.
**$p < 0.05$.
*** $p < 0.01$.

were considered. Focussing on the effects of the rate of autonomous improvement, α, we see that no statistically significant effect on the fatalities objective is observed. However, the effect on the suffering objective is judged to be significant ($p < 0.01$). While statistically significant, it should be emphasised that the effect size itself is modest. Estimated to be 1169, this corresponds to an change in suffering of 117 units for any 0.1 change in the rate of autonomous improvement.

5.4. Discussion

The routing strategies described in Section 4 were noted to have different expected effects on both the length of the resulting routes and the predictability of the associated travel times. The strategy of static routing (SR) was noted to favour the ability to accurately predict travel times at the expense of finding shorter routes. This is achieved through consistently using a single route for each journey, namely that which would be expected to be the shortest using baseline data describing the transport network. By using the same route for each journey, the optimisation model is capable of revising its estimate of the travel time on that route as the relevant data is collected, leading to improved predictions as the response operation progresses. The strategies of CAR, AIAR and ACAR, on the other hand, sacrifice this ability to 'learn' the travel times of a specific route, as the routes taken for each journey are allowed to change each time in the hope of finding shorter routes.

The analysis presented in this subsection demonstrates that the choice of routing strategy does not have any significant impact upon the utility of the optimisation model when assessing performance through the fatalities objective. By this measure, the strategies of SR, CAR and ACAR all lead to similar performance. In terms of the objective of suffering, however, some differences can be noted. In particular, autonomous routing can lead to improved performance in certain situations. Specifically, in scenarios where the level of disruption to the network is large, where responders are assumed able to share knowledge regarding routing (that is, where the routing strategy of ACAR as opposed to AIAR is adopted), and where the rate of autonomous improvement is set to $\alpha = 0.8$, autonomous routing can lead to improved performance. This result is of interest, particularly when viewed in the context of related research into decision support for large scale emergency response involving routing decisions. As was discussed in the review of such related work given in Section 1, it is common for decision support to include the specification of which route responders should take when enacting response operations.

6. Conclusions and further work

The analysis presented in the paper has explored the effect of routing on performance whilst using an optimisation model in MCI response. Specifically, we have described four policies defining how routing decisions should be made in an MCI and evaluated, through Monte Carlo experiments, their utility. In doing so, we not only considered the ability of each policy to find routes with low travel time, but also their consistency, that is, their ability to produce routes with predictable travel times. This predictability is of paramount importance when employing an optimisation model which relies on an ability to forecast the future implications of current decisions, such as the scheduling model used in this paper.

In this paper we have shown that, if the proposed scheduling based optimisation model is to be used in MCI response, it may be beneficial to remove routing decisions from the programs remit and leave these to the responders to determine themselves, providing they are able to share knowledge and learn together. Moreover,

the relationship between the utility gained through this approach and the level of disruption in the transport network has been quantified. Finally, we have described an improved methodology for where centralised routing decisions are required, employing a Bayesian approach to updating beliefs regarding the distributions of travel times associated with journeys.

6.1. Further work

One central limitation of the analysis described, particularly for static routing (*SR*), is the assumption of a constant underlying disruption parameter on each link. Although we can argue this is valid for short response operations (say, around 1 or 2 h), these parameters will likely vary over time during longer incidents. The other policies presented should adapt well to this scenario, although further empirical investigations are required to confirm that this is the case. Another potentially limiting factor of the disruption model is that each link's parameter is uncorrelated. In reality, we would expect a high level of correlation between parameters of adjacent links. If this were to be modelled the performance of the centralised adaptive policy (*CAR*) could be improved by updating the whole network in a manner which reflects the covariance structure. This would not be a trivial task, however, and could require unrealistic or undesirable computational resources.

In practical terms, this study could be greatly improved if data were available on routing and travel times in major incident response. This remains unlikely due to the inherently low frequency of such events. More generally, insights should be extracted from this work regarding temporal variations, dynamic information flow and control in the response environment. Future work will examine these factors in a wider sense, considering what problem and response characteristics lead to dynamic behaviour and uncertainty, and how these factors in turn influence the utility of the decision support program.

Acknowledgements

References

Advanced Life Support Group, 2011. Major Incident Medical Management and Support: The Practical Approach at the Scene, third ed. Wiley-Blackwell.

Altaya, N., Green, W.G., 2006. OR/MS research in disaster operations management. Eur. J. Oper. Res. 175 (1), 475–493. http://dx.doi.org/10.1016/j.ejor.2005.05.016.

Barbarosoglu, G., Arda, Y., 2004. A two-stage stochastic programming framework for transportation planning in disaster response. J. Oper. Res. Soc. 55 (1), 43–53. http://dx.doi.org/10.1057/palgrave.jors.2601652.

Brandimarte, P., 1993. Routing and scheduling in a flexible job shop by tabu search. Ann. Oper. Res. 41, 157–183.

Budge, S., Ingolfsson, A., Zerom, D., 2010. Empirical analysis of ambulance travel times: the case of Calgary emergency medical services. Manage. Sci. 56, 716–723. http://dx.doi.org/10.1287/mnsc.1090.1142.

Caunhye, A.M., Nie, X., Pokharel, S., 2012. Optimization models in emergency logistics: a literature review. Socio-Econ. Plan. Sci. 46, 4–13. http://dx.doi.org/10.1016/j.seps.2011.04.004.

Chang, M.-S., Tseng, Y.-L., Chen, J.-W., 2007. A scenario planning approach for the flood emergency logistics preparation problem under uncertainty. Transport. Res. Part E: Logist. Transport. Rev. 43 (6), 737–754. http://dx.doi.org/10.1016/j.tre.2006.10.013.

Chiu, Y.-C., Zheng, H., 2007. Real-time mobilization decisions for multi-priority emergency response resources and evacuation groups: model formulation and solution. Transport. Res. Part E: Logist. Transport. Rev. 43 (6), 710–736. http://dx.doi.org/10.1016/j.tre.2006.11.006.

de la Torre, L.E., Dolinskaya, I.S., Smilowitz, K.R., 2012. Disaster relief routing: integrating research and practice. Socio-Econ. Plan. Sci. 46, 88–97. http://dx.doi.org/10.1016/j.seps.2011.06.001.

Fiedrich, F., Gehbauer, F., Rickers, U., 2000. Optimized resource allocation for emergency response after earthquake disasters. Safety Sci. 35, 41–57.

Gong, Q., Batta, R., 2007. Allocation and reallocation of ambulances to casualty clusters in a disaster relief operation. IIE Trans. 39 (1), 27–39. http://dx.doi.org/10.1080/07408170600743938.

Haghani, A., Oh, C., 1996. Formulation and solution of a multi-commodity, multi-modal network flow model for disaster relief operations. Transport. Res., Part A 30 (3), 231–250.

Hawe, G.I., Coates, G., Wilson, D.T., Crouch, R.S., 2012. Investigating the effect of overtriage on hospital arrival times of critically injured casualties during a major incident using agent-based simulation. In: Proceedings of the 6th International Conference on Soft Computing and Intelligent Systems and the 13th International Symposium on Advanced Intelligent Systems (SCIS-ISIS).

Jiang, Y., Yuan, Y., Huang, K., Zhao, L., 2012. Logistics for large-scale disaster response: achievements and challenges. In: 45th Hawaii International Conference on System Sciences. http://dx.doi.org/10.1109/HICSS.2012.418.

Jotshi, A., Gong, Q., Batta, R., 2009. Dispatching and routing of emergency vehicles in disaster mitigation using data fusion. Socio-Econ. Plan. Sci. 43, 1–24. http://dx.doi.org/10.1016/j.seps.2008.02.005.

Kolesar, P., Walker, W., Hausner, J., 1975. Determining the relation between fire engine travel times and travel distances in New York City. Oper. Res. 23, 614–627. http://dx.doi.org/10.1287/opre.23.4.614.

Lin, Y.-H., Batta, R., Rogerson, P.A., Blatt, A., Flanigan, M., 2011. A logistics model for emergency supply of critical items in the aftermath of a disaster. Socio-Econ. Plan. Sci. 45, 132–145.

Mete, H.O., Zabinsky, Z.B., 2010. Stochastic optimization of medical supply location and distribution in disaster management. Int. J. Prod. Econ. 126, 76–84. http://dx.doi.org/10.1016/j.ijpe.2009.10.004.

Ozdamar, L., Ekinci, E., Kucukyazici, B., 2004. Emergency logistics planning in natural disasters. Ann. Oper. Res. 129, 217–245.

Rawls, C.G., Turnquist, M.A., 2010. Pre-positioning of emergency supplies for disaster response. Transport. Res. Part B: Methodol. 44, 521–534. http://dx.doi.org/10.1016/j.trb.2009.08.003.

Rolland, E., Patterson, R.A., Ward, K., Dodin, B., 2010. Decision support for disaster management. Oper. Manage. Res. 3, 68–79. http://dx.doi.org/10.1007/s12063-010-0028-0.

Sheu, J.-B., 2007. An emergency logistics distribution approach for quick response to urgent relief demand in disasters. Transport. Res. Part E: Logist. Transport. Rev. 43, 687–709. http://dx.doi.org/10.1016/j.tre.2006.04.004.

Simpson, N., Hancock, P., 2009. Fifty years of operational research and emergency response. J. Oper. Res. Soc. 60, SI26–SI39. http://dx.doi.org/10.1057/jors.2009.3.

Skiena, S., 1990. Dijkstra's Algorithm, Implementing Discrete Mathematics: Combinatorics and Graph Theory with Mathematica. Addison-Wesley, Reading, MA, pp. 225–227.

Tzeng, G.-H., Cheng, H.-J., Huang, T.D., 2007. Multi-objective optimal planning for designing relief delivery systems. Transport. Res. Part E: Logist. Transport. Rev. 43, 673–686. http://dx.doi.org/10.1016/j.tre.2006.10.012.

Westgate, B.S., Woodward, D.B., Matteson, D.S., Henderson, S.G., 2011. Travel Time Estimation for Ambulances Using Bayesian Data Augmentation, Tech. Rep., Cornell University. <http://hdl.handle.net/1813/23135>.

Wex, F., Schryen, G., Neumann, D., 2012. Operational emergency response under informational uncertainty: a fuzzy optimization model for scheduling and allocating rescue units. In: Proceedings of the 9th International ISCRAM Conference.

Wilson, D.T., Hawe, G.I., Coates, G., Crouch, R.S., 2012. Stochastic task durations in the scheduling of mass casualty incident response operations. In: Proceedings of the 30th Workshop of the UK Planning and Scheduling Special Interest Group (PlanSIG).

Wilson, D.T., Hawe, G.I., Coates, G., Crouch, R.S., 2013. A multi-objective combinatorial model of casualty processing in major incident response. Eur. J. Oper. Res. 230 (3), 643–655. http://dx.doi.org/10.1016/j.ejor.2013.04.040.

Wilson, D.T., Hawe, G.I., Coates, G., Crouch, R.S., 2013. Scheduling response operations under transport network disruptions. In: Proceedings of the 10th International ISCRAM Conference.

Yi, W., Kumar, A., 2007. Ant colony optimization for disaster relief operations. Transport. Res. Part E: Logist. Transport. Rev. 43, 660–672. http://dx.doi.org/10.1016/j.tre.2006.05.004.

Yi, W., Ozdamar, L., 2007. A dynamic logistics coordination model for evacuation and support in disaster response activities. Eur. J. Oper. Res. 179, 1177–1193. http://dx.doi.org/10.1016/j.ejor.2005.03.077.

Yuan, Y., Wang, D., 2009. Path selection model and algorithm for emergency logistics management. Comput. Ind. Eng. 56, 1081–1094. http://dx.doi.org/10.1016/j.cie.2008.09.033.

Zhang, J.-H., Li, J., Liu, Z.-P., 2012. Multiple-resource and multiple-depot emergency response problem considering secondary disasters. Expert Syst. Appl. 39, 11066–11071.

Zhang, X., Zhang, Z., Zhang, Y., Wei, D., Deng, Y., 2013. Route selection for emergency logistics management: a bio-inspired algorithm. Safety Sci. 54, 87–91. http://dx.doi.org/10.1016/j.ssci.2012.12.003.

Zhou, Q., Huang, W., Zhang, Y., 2011. Identifying critical success factors in emergency management using a fuzzy DEMATEL method. Safety Sci. 49, 243–252, http://dx.doi.org/10.1016/j.ssci.2010.08.005.