School of Mathematics, Cardiff University, Abacws, CF24 4AG, Cardiff, Wales, United Kingdom.

Welsh Ambulance Services NHS Trust, Beacon House, NP44 3AB, Cwmbran, Wales, United Kingdom

Aneurin Bevan University Health Board, Lodge Road, NP18 3XQ, Caerleon, Wales, United Kingdom

s://doi.org/10.1093/imaman/dpaf027

On the Optimisation of Heterogeneous Ambulance Fleet Allocations

GERAINT IAN PALMER*, MARK TUSON, VINCENT KNIGHT, PAUL HARPER, AND SARIE BRIEE
ichool of Mathematics, Cardiff University, Abacws, CF24 4AG, Cardiff, Wales, United Kingdom
LEANNE SMITH
Ish Ambulance Services NHS Trust, Beacon House, NP44 3AB, Cwmbran, Wales, United Kingdom
COTESPONDING author: palmergil @Cardiff.ac.bk

GRecaived on 24 January 2024; accepted on 1 July 2025]

Ambulance services have a duty of care to the clinical outcomes of the population they serve, and
terrefore aim to maximise the chances of survival and improve patient outcomes following a medical
emergency. Despite this, although the ambulance allocation problem has been widely studied, it has predominantly focused on minimising response times or maximising coverage alone, and not explicitly for
considering patient outcomes. In this paper we profose a pidelling approach to consider where to best
allocate different types of emergency response vehicles. Traditionally this utilisation is problematic as it is both an input and an output of the allocation of vehicles. Our approach is informed by,
and tested on, real-world data from-fakata, Indonesia. Using our developed models, decision makers
are better able to understand ambulance fleet capacity needs and allocations, and their impact on patient
outcomes.

Pywords: Ambulance Planning: Emergency Medical Services; Simulation; Health Care; Optimisation;
Pharematical Modelling
The Author(s) 2025. Published by Oxford University Press. This is an Open Access artice
thrist.//creativecommons.org/licenses/by/4.0/), which permits unrestricted reuse, distribution
with a barefore the version of the creative Commons Attribution License
thrist.//creativecommons.org/licenses/by/4.0/), which permits urrestricted reuse, distribution

distributed under the terms of the Creative Commons Attribution License (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted reuse, distribution and reproduction in any medium, provided the original work is properly cited.

1. Introduction and Background

Time-critical conditions (TCCs) are a substantial cause of mortality worldwide, responsible for an estimated 54% of all deaths (Fraser et al., 2020). Emergency medical services (EMS) play a pivotal role in responding to TCCs and saving patients from life-threatening medical conditions. Most EMS research, especially planning to meet increasing ambulance demands, tends to be focussed on high-income countries including, for example, Australia (Lowthian et al., 2011), Germany (Veser et al., 2020), and USA (Birmingham et al., 2021). In contrast, many low and middle-income countries (LMICs) lack an organised EMS system, with most ambulances used purely for patient transport and not as an emergency care vehicle (Suryanto et al., 2017). However, the burden of deaths due to TCCs is much greater in LMICs than in high-income countries, the difference being around threefold (Chang et al., 2016). For traffic accidents, strokes and heart attacks, rapid access to appropriate treatment is especially vital.

The application of our research is focused on Indonesia, supported by an award for Global Challenges Research Funds (GCRF). In Indonesia, an LMIC, there has been very little prior work on developing an EMS strategy, driven by a lack of financial investment and human capital (Pusponegoro, 2003, Suryanto et al., 2017, Yusvirazi et al., 2018). Related research by the authors (Brice et al., 2022) has highlighted particular challenges in Indonesia as a lack of a single coordinated EMS, the ability or willingness to pay for an ambulance, a large geographical area, and areas of severe traffic congestion especially in Jakarta, the capital.

When our research programme commenced in October 2019, ambulance services in Jakarta were provided by many disparate, mostly private providers that charge patients for their use. We initially partnered with Ambulans 118, a non-government charitable ambulance service established in 2005 by the Indonesian Surgeons Association, that currently operates in five cities across Indonesia: Jakarta, Palembang, Yogyakarta, Surabaya and Makasser. Unlike private providers, Ambulans 118 suggests that a donation is made for use of its EMS vehicles, but otherwise provides free emergency medical care, it is also leading on paramedic training across the country.

The overarching goal of our research was to work with the Indonesian Government, the different ambulance providers, and hospitals, to help them forecast emergency demand and make critical decisions on the best types, capacities and geographical allocations of emergency vehicles within a potentially co-ordinated EMS system, starting with Jakarta. To facilitate this, in collaboration with Ambulans 118 we organised workshops in Jakarta in February 2020, which were attended by over 170 people including doctors, nurses, paramedics, academics, and officials from the Indonesian Ministry of Health. This paper presents some aspects of the overall research programme, specifically focusing on the development of a simulation model and allocation metaheuristic for maximising patient survival.

Fundamental to the project and the research described in this paper was an initial study that we carried out in order to better understand patient needs and the barriers to use of ambulances in Indonesia. Throughout the month of December 2020, we undertook comprehensive surveys in Emergency Departments (EDs) across Jakarta and published the first known study in to EMS demand within the country (Brice et al., 2022). Our study showed that the utilisation of ambulances by patients attending EDs is very low and of concern. The low utilisation is contributed by patients' lack of awareness of available ambulance services, patients' disinclination to use ambulances due to high costs, and long response times. All of these barriers impact on patient outcomes, especially for those with life-threatening conditions. For example, more trauma patients took a car-share ride (20%) or motorcycle (20%) to reach the ED than an ambulance (just 10%), while only 14% of critical cardiovascular patients used an ambulance compared to 67% travelling to hospital by private car or a car-share ride.

The first contribution of this paper is an extension the work of Knight et al. (2012) to consider both heterogeneous patients and heterogeneous fleets. Here a maximal survival objective is formulated, from vehicle allocations and utilisations, which are themselves found by numerically solving a relationship on the vehicles' share of the demand. This is then used in a metaheuristic to find improved allocations that aim to maximise patient survival. Secondly, we develop a comprehensive sequential discrete event EMS simulation that models and evaluates heterogeneous fleet allocations of emergency vehicles by utilising a novel approach in which transit jobs, rather than patients, are framed as the queueing 'customers'. Our methodology utilises two discrete event simulations of the same system that are run sequentially and together combine to form the logic of a single simulation of a heterogeneous fleet; and key performance indicators (KPIs) are calculated by making use of survival functions. Thirdly, we demonstrate the use and impact of our modelling framework applied to current and proposed ambulance allocations in the city of Jakarta, thus supporting Government-level decision making with an overall goal to improve the lives of those living in LMICs. This enables decision support and managerial insights in Indonesia, although the developed modelling framework could be readily applied to other locations.

There are several managerial implications for EMS providers and government agencies in Jakarta and Indonesia, respectively. The findings emphasise the importance of adopting a dynamic and data-driven approach to resource allocation. The inclusion of a heterogeneous fleet, such as Rapid Response Vehicles (RRVs) alongside traditional ambulances, has been shown to significantly improve response times and patient survival rates, particularly in densely populated urban areas. Managers are advised to invest in these diverse fleets and adopt flexible allocation strategies that account for temporal variations in demand, such as the daily influx of commuters and the differing population densities across neighbourhoods. The research underscores the necessity of enhancing data collection practices to support informed decision-making, with the development of user-friendly decision support tools to assist in visualising data and predicting outcomes.

The paper is structured as follows: Section 2 gives an overview of the literature on modelling ambulance allocations. Section 3 describes the current ambulance service behaviour and sets out the mathematical notation used throughout the paper. Section 4 describes an optimisation model, including the survival objective, utilisation considerations, and a metaheuristic algorithm to find better fleet allocations. Section 5 outlines the logic of the sequential discrete

event simulation models for heterogeneous fleets. Section 6 gives a case study, describing the current emergency service situation in Jakarta and applying the optimisation and simulation results. Section 7 discusses the findings and contributions of the paper.

2. Literature Review

There is a rich literature on operational research applied to EMS location and related allocation problems. Literature reviews that include these topics are Aringhieri et al. (2017), Bélanger et al. (2019), Farahani et al. (2019), Li et al. (2011), Liu et al. (2021), Mukhopadhyay et al. (2022), Reuter-Oppermann and Vile (2017) and Wang et al. (2021). Ambulance allocation problems such as that concerned in this paper can be categorised as a *preparedness* problem (Mukhopadhyay et al., 2022).

Emergency medical services can be evaluated using different performance measures with coverage and response time being the predominant metrics (McLay and Mayorga, 2010). However, missing a response time target by just one second, for example, would be considered as a 'failure' in many models with no appreciation at all of the impact on patient survival or outcome (McLay and Mayorga, 2010). This seems somewhat restrictive and short-sighted, which is why we focus on survival as a metric and compare and contrast papers using that metric with our approach. Here we consider some closely related papers, providing a summary of each and comparing and contrasting them with our modelling and solution approach.

Amorim et al. (2019) propose an integrated strategic and tactical planning approach which features an optimisation model and a local search heuristic based on Gaussian Processes. Their methodology is applied to the city of Porto while reporting on performance metrics such as survival. Our approach is different because we use a discrete-event simulation framework to evaluate results from a population based heuristic to determine the number of ambulances required in each station.

Boutilier and Chan (2022) develop an integrated location-queueing model that incorporates existing EMS response times in a drone network. They use a p-median approach and an Erlang loss model. Although survival is not their main optimisation criterion, they report how many patients would survive using their approach.

Erkut et al. (2008) can be considered the first modelling approach that considers decaying survival probabilities during response time. The rationale is, because survival can be thought of as a more robust and generic objective for EMS performance measurement than coverage or average response time. In an application of a recently developed Maximal Survival Location Problem model (MSLP), the authors use data from Edmonton, Canada and show that maximising survival is superior to other objectives in clinical outcome for cardiac arrest patients.

McCormack and Coates (2015) focus on the optimisation of EMS vehicle fleet allocation and base station location through the use of a genetic algorithm (GA) with an integrated EMS simulation model. Their objective is maximization of the overall expected survival probability across patient classes. Applications of the model were undertaken using real call data from the London Ambulance Service. The difference between their modelling approach and our is that we have a different survival function, focus on a different heuristic optimisation and evaluate our approach with data from a developing country, and have higher traffic volume and congestion in our data.

Bélanger et al. (2020) combine optimisation and simulation to find both ambulance locations and dispatch policies. Here, binary integer programming is used to minimise total response time, where ambulance availability, or utilisation is unknown. In this work this is overcome by iteratively feeding the solution of the optimisation programme into a simulation to find the availability parameters, until convergence. Our work differs as we solve for ambulance utilisation explicitly.

Toro-Díaz et al. (2013, 2015) build optimisation models for both the allocation and the dispatch, or preference, rules for EMS services. They utilise the hypercube queueing model for ambulance availability with a integer programming and metaheuristic algorithms. In particular, they consider a number of different objectives including mean response times, expected coverage (Daskin, 1983), and the Gini index on individual response times as measures of fairness in the system.

Knight et al. (2012) provide a supporting extension of the MSLP that allows it to be applied more generally to real-world EMS systems, acknowledging that different patients have varying levels of expected survival probabilities. The Maximum Expected Survival Location Model for Heterogeneous Patients, MESLMHP, allows multiple classes of patient groups to be defined, where previously only cardiac arrest patients formed part of the objectives for primary response. In reality any emergency patient has a necessity for swift attendance, and a timely response to any incident type may impact on clinical outcome in some way. It is for this reason MESLMHP is designed to be generic enough to accommodate any number of patient groups, with each class weighted dependent upon the relative urgency of the incident. More than just a contribution to the models demand input, these patients are included in the optimisation when maximising total population survival probability. Our work directly extends this work, to include heterogeneous vehicle fleets, while developing a different fixed-point numerical solution to the problem of approximating vehicle utilisations.

3. Problem Statement & Notation

We first present some notation. We have the following sets:

- \mathscr{P} is the set of pick-up locations, indexed by p;
- \mathscr{A} is the set of ambulance locations, indexed by a;
- \mathcal{Y} is the set of hospitals, indexed by y;
- \mathcal{K} is the set of medical specialities, indexed by k;

to incorporate heterogeneous fleets the set of specialities \mathcal{H} is partitioned into two sets, \mathcal{H}_A , those patients that can be seen by secondary vehicles, and \mathcal{H}_B , those patients who are seen by primary vehicles only.

An allocation is given by the tuple (Z_a, \tilde{Z}_a) , where Z_a is the number of primary vehicles allocated to ambulance location a; and \tilde{Z}_a is the number of secondary vehicles allocated to ambulance location a. Also let:

- λ_{pk} be the rate at which patients of speciality k make calls from pick-up location p;
- $s_k(t)$ be the survival function function associated with patients of speciality k;
- π_a be the average utilisation of primary vehicles stationed at location *a*;
- b_{pa} be the expected travel time from ambulance location *a* to pick-up location *p*; and let B_{pa} be a random variable representing the time it takes for an ambulance to drive this distance;
- C_{py} be a random variable representing the travel time from pick-up location p to hospital y;
- *D_{ya}* be a random variable representing the travel time from hospital *y* to ambulance location *a*;
- *F_{pa}* be a random variable representing the travel time from pick-up location *p* to ambulance location *a*;
- *G_k* be a random variable representing the time the ambulance spends with patients of speciality *k* at the pick-up location;
- J_k be a random variable representing the time the ambulance spends with patients of speciality k at the hospital;
- Θ be a random variable representing the time the ambulance spends re-fuelling, re-stocking, and resting between transit jobs;
- q_{pky} be the probability that a patient of speciality k from pick-up location p is taken to hospital y. Note that having $\sum_{y \in \mathscr{U}} q_{pky} < 1$ is possible, that is a patient may not go to any hospital, in which case the ambulance returns to their ambulance location.

To incorporate heterogeneous vehicles, let also $\tilde{\pi}_a$ denote the average utilisation of secondary vehicles stationed at location *a*; and \tilde{b}_{pa} , \tilde{B}_{pa} , \tilde{C}_{py} , \tilde{D}_{ya} , and \tilde{F}_{pa} denote the corresponding travel times for secondary vehicles.

Patients will call an ambulance from one of the pick-up locations $p \in \mathcal{P}$. If *all* ambulances are busy, then that call will be abandoned and it is assumed that the patient will find their own way to the hospital through private or public transportation, which is an appropriate and justified assumption for the case of Jakarta (Brice et al., 2022). Otherwise, a central control centre will dispatch the closest available primary vehicle. That vehicle will travel from its location to the patient (B_{pa}), spend some time on the scene treating the patient (G_k), travel to the hospital (C_{py}), spend some time handing over the patient at the hospital (J_k), travel back to its original location (D_{ya}), then spend time refilling and refuelling (Θ). There is also a probability that the patient does not need a hospital, and so the vehicle will travel from the pick-up location back to it's original vehicle location (F_{pa}). These routes are shown visually in Figure 1. Note that these routes cannot be interrupted, that is, an ambulance must return to



FIG. 1. Ambulance routes for a given transit job.

it's station for refilling before it is available to respond to another call, to ensure it is cleaned, and fully stocked for it's next call. Note also that B_{pa} and F_{pa} need not necessarily be equal, as travel times need not necessarily be symmetrical, this could be due to a number of reasons, such as non-symmetric road networks, and differences in driver urgency on each leg of the journey.

Therefore the total time an ambulance is busy will be T_{paky} given by Equation 1 if a hospital is required, and service time T_{pak} given by Equation 2 if no hospital is required.

$$T_{paky} = B_{pa} + G_k + C_{py} + J_k + D_{ya} + \Theta$$
⁽¹⁾

$$T_{pak} = B_{pa} + G_k + F_{pa} + \Theta \tag{2}$$

There are two types of vehicle: primary vehicles, typically ambulances, that must be dispatched to all patients; and secondary vehicles, typically rapid response motorcycles, that can travel faster than primary vehicles, and are dispatched to some patients to respond to an emergency faster than the primary vehicle. Secondary vehicles cannot transport patients, but are used in conjunction with primary vehicles to reduce response times. From a resource cost perspective, one primary vehicle costs the same to purchase and to run as three secondary vehicles.

4. Optimising Allocations for Maximal Survival

Here we outline a method of finding vehicle allocations that maximise the expected survival of patients. The problem is, for a given set of vehicle locations (that is ambulance stations), a given total number of primary vehicles, and a given total number of secondary vehicles, how many vehicles of each type should be stationed at each location to maximise expected survival across all patients.



FIG. 2. Survival function s(t), estimated by Valenzuela et al. (2000) compared with $s_L(t)$ with a hard cut-off of L = 8 represented as a step function.

4.1. Survival Functions

A key concept here is the survival curve of a patient. In reality, some emergency incidents do not result in a substantive deterioration of a patients status over time; however, in all situations, there is a reasonable cut-off beyond which a patient should not expect to have to wait for care, and mixture of theoretical survival functions and step functions can be utilised. Survival probabilities for critical incidents, calculated from a theoretical monotonically decaying survival function reported in the literature (Valenzuela et al., 2000) are used to demonstrate an attainable level of success from a response. One particular survival curve s(t) of Equation 3, represents survival until hospital discharge following cardiac arrest; its origins are explained in detail by Knight et al. (2012), and gives the probability of survival if seen within a time t in the form of a logistic function. Figure 2 shows the difference between using this survival curve and a hard cut-off of 8 minutes. However, hard cut-off curves like Equation 4, with a cut-off of L can still be used to represent meeting artificially selected targets, for example for transportation jobs.

$$s(t) = \left(1 + e^{0.26 + 0.139t}\right)^{-1} \tag{3}$$

$$s_L(t) = \begin{cases} 1 & \text{if } 0 \le t \le L \\ 0 & \text{if } t > L \end{cases}$$

$$\tag{4}$$

Finding allocations that maximise survival can result in allocations with different behaviours to those found by minimising response times or maximising coverage. Coverage maximisation models are often easier to solve, and easier to communicate, but may not take into consideration severity or intensity of demand (Erkut et al., 2008). This results in a minority of patients, particularly those more isolated, being over-serviced, that is, having more relative resources allocated to their needs that the majority, at the expense of the majority. Minimising average response times on the other hand treats all patients equally, but may have the unintended outcome of disregarding some more further out patients whose response times would be larger that others regardless of the allocation. It would still disregard the severity or type of the patient. Allocations based on maximising survival are similar to those that minimise response time, overcoming the issue of over-servicing outlying patients, though it is still possible to disregard them completely. It has an advantage over minimising response times as it can account for patient severity. Having survival as functions of response times has the effect of weighting gains in response times by how valuable those gains are to a patient.

4.2. Maximal Expected Survival Location Model with Heterogeneous Patients and Heterogeneous Fleet

Here we propose a Maximal Expected Survival Location Model with Heterogeneous Patients and Heterogeneous Fleet (MESLMHPHF). It is an extension of the MESLMHP model given in Knight et al. (2012), which did not consider heterogeneous fleets. This is a model of survival that can be used as an objective function for optimisation algorithms. It is a weighted expected survival function as the objective function which considers heterogeneous patients (different specialities) and heterogeneous fleets (both primary vehicles, EAs, and secondary vehicles, RRVs), and is constructed by appropriately summing these survival curves, described in Section 4.1, across the population, multiplying by ambulance availability where appropriate. The full MESLMHPHF optimisation model is given in APPENDIX B, and its components are described fully here. The MESLMHPHF is given by Equation 5.

$$g\left(Z_{a},\tilde{Z}_{a}\right) = \sum_{p \in \mathscr{P}} \sum_{a \in \mathscr{A}} \left(\sum_{k \in \mathscr{K}_{A}} w_{k} \lambda_{pk} \hat{\Psi}_{kpa} + \sum_{k \in \mathscr{K}_{B}} w_{k} \lambda_{pk} \Psi_{kpa} \right)$$
(5)

where w_k is a weight associated with patient speciality type k. For this study we assume $w_k = 1$ for all k, which allows the interpretation of $g(Z_a, \tilde{Z}_a)$ as the expected number of patients surviving per time unit. Now Ψ_{kpa} can be interpreted as the probability of a patient $k \in \mathcal{K}_B$ at pick-up location p being seen by a primary vehicle from location a and surviving; while Ψ_{kpa} would be the probability of patients of speciality $k \in \mathcal{K}_A$ at pick-up location p being seen by any vehicle from location z and surviving. Equation 5 is the weighted sum over the expected survival probabilities of all patient specialities, all patient pick-up locations, and all ambulance stations. These survival probabilities are given by Equations 6 and 7 respectively.

$$\Psi_{kpa} = s_k \left(b_{pa} \right) \left(1 - \pi_a^{Z_a} \right) \prod_{\alpha \in \mathscr{A}} \pi_\alpha^{\left(Z_\alpha \beta_{p\alpha a} \right)} \tag{6}$$

$$\begin{split} \hat{\Psi}_{kpa} &= s_k \left(\tilde{b}_{pa} \right) \left(1 - \tilde{\pi}_a^{\tilde{Z}_a} \right) \prod_{\alpha \in \mathscr{A}} \tilde{\pi}_{\alpha}^{\left(\tilde{Z}_{\alpha} \beta_{p\alpha a} \right)} \pi_{\alpha}^{\left(Z_{\alpha} R_{p\alpha a} \right)} \\ &+ s_k \left(b_{pa} \right) \left(1 - \pi_a^{Z_a} \right) \prod_{\alpha \in \mathscr{A}} \pi_{\alpha}^{\left(Z_{\alpha} \beta_{p\alpha a} \right)} \tilde{\pi}_{\alpha}^{\left(\tilde{Z}_{\alpha} \left(1 - R_{pa\alpha} \right) \right)} \end{split}$$

Here $R_{pa_1a_2}$ and $\beta_{pa_1a_2}$ define an ordering of the ambulance locations, they are binary variables indicating the preference of sending a vehicle from one station to another, that is the dispatch rule. Here we use the closest or fastest vehicle, so $\beta_{pa_1a_2}$ indicates if a vehicle of the same type can reach p quicker from a_1 than a_2 , defined in Equation 8, while $R_{pa_1a_2}$ indicates if a primary vehicle at a_1 can reach p quicker than a secondary vehicle at a_2 , defined in Equation 9. In order for these to define an ordering it is assumed that no two distinct ambulance locations are exactly equidistant from a pickup location. In both Equation 8 and 9, \leq are used rather than <; in Equation 8 will be equivalent, as we assume no two ambulance locations are equidistant; while in Equation 9, in the case of a primary and secondary vehicle having equal travel times, then the primary vehicle will have priority. Note that here these parameters are defined using travel times, but can be generalised to account for any type of preference, including for example proximity, importance, or efficiency.

$$\beta_{pa_{1}a_{2}} = \begin{cases} 0 & \text{if } a_{1} = a_{2} \\ 1 & \text{if } b_{pa_{1}} \le b_{pa_{2}} \\ 0 & \text{otherwise.} \end{cases}$$

$$R_{pa_{1}a_{2}} = \begin{cases} 1 & \text{if } b_{pa_{1}} \le \tilde{b}_{pa_{2}} \\ \end{cases}$$
(8)
(9)

the equations in APPENDIX A.

Note that in Equations 6 and 7, the survival functions used, $s_k(t)$, are the same whether a primary or a secondary vehicle arrives first. This may be unrealistic, where primary vehicles may be better equipped than secondary vehicles to deal with some emergencies. One option to overcome this would be to have distinct survival functions for each vehicle type, however there may be some interactions between the survival probability functions that would not be captured, for example when a primary vehicle arrives after a secondary vehicle, there may still be a better chance of survival that the secondary vehicle only. This would be an interesting

(7)

avenue of further research. For this work, survival functions will be the same for all vehicles, representing the probability of surviving given time to be seen by any emergency vehicle. Note also that the survival functions are functions of b_{pa} , the expected travel time between p and a, and do not consider the variability in the travel times. This will be discussed further alongside results in Section 6.4.

A key consideration is the vehicle utilisations π_a and $\tilde{\pi}_a$; discussed in the next subsection.

4.3. Utilisation Considerations

The models above assume that the utilisations of each vehicle type at each vehicle location is known, which is a potentially restrictive assumption. These utilisations π_a and $\tilde{\pi}_a$ depend on the demand to station *a*, which itself depends on the vehicle allocations. One method of overcoming this, used in Knight et al. (2012), is to consider utilisation as the ratio of demand and service rates, shown in Equations 10 and 11. Here we let $\frac{1}{\mu}$ and $\frac{1}{\tilde{\mu}}$ be the average job times of primary and secondary vehicles, then λ_a and $\tilde{\lambda}_a$ represent the share of the demand seen by primary and secondary vehicles from vehicle location *a*, respectively.

$$\pi_{a} = \frac{\lambda_{a}}{\mu}$$
(10)
$$\tilde{\pi}_{a} = \frac{\tilde{\lambda}_{a}}{\tilde{\mu}}$$
(11)

Note that here we assume that average job times are not dependent on the vehicle location a, although this may be unrealistic, given that the travel between locations are a key part of a vehicle's job time (B_{pa}, D_{ya}, F_{pa}) . However, as discussed in Section 3, there are a number of other components to an ambulance job including time on site (G_k) , hospital handover time (J_k) , and vehicle refill time (Θ) . For the case of Jakarta, we justify the assumption of job times not being location dependent by considering the proportion of an ambulance job's time that is location dependent, ρ , defined in Equation 12,

$$\rho = \begin{cases} \frac{B_{pa} + D_{ya}}{T_{paky}} & \text{if hospital required,} \\ \frac{B_{pa} + F_{pa}}{T_{pak}} & \text{if no hospital required,} \end{cases}$$
(12)

then, using the simulation described in Section 5, we consider the distribution of ρ across all simulated jobs under the current situation and allocation in Jakarta, shown in Figure 3. We see for primary vehicle jobs on average 24% of a job time is location dependent, while it is 36.5% for secondary vehicles. Considering location-dependent service rates would be an interesting further research question, potentially improving the model's performance by incorporating travel times into the MESLMHPHF. However, given that in the case of Jakarta the preferences $\beta_{pa_1a_2}$ and $R_{pa_1a_2}$ are themselves defined by travel times (Equations 8 and 9), this might not

12 ON THE OPTIMISATION OF HETEROGENEOUS AMBULANCE FLEET ALLOCATIONS



FIG. 3. Distribution of the proportion of an ambulance job consists of location dependent components, for primary and secondary vehicles.

have a strong impact on the results of may not have much affect for this study, as the model already gives preference to the locations with the shortest service times.

The demand experienced by vehicles at each location $a \in \mathscr{A}$ would indeed be location specific, and would depend on the number and utilisation of the vehicles at every other location. This is addressed in the next subsection.

4.4. Numerical approximation of utilisation

For primary vehicles, which operate independently of secondary vehicles, the relationship between demand experienced at each location, λ_a , and the allocations and utilisations of all other locations, is given by Equation 13. For secondary vehicles, the relationship between demand experienced at each location, $\tilde{\lambda}_a$, and the allocations and utilisations of all other locations, is given by Equation 14. Note that this also depends on the utilisations of the primary vehicles, π_a .

$$\lambda_{a} = \sum_{p \in \mathscr{P}} \sum_{k \in \mathscr{K}} \lambda_{pk} \left(1 - \left(\frac{\lambda_{a}}{\mu}\right)^{Z_{a}} \right) \prod_{\alpha \in \mathscr{A}} \left(\frac{\lambda_{\alpha}}{\mu}\right)^{\left(Z_{\alpha}\beta_{p\alpha a}\right)}$$
(13)

$$\tilde{\lambda}_{a} = \sum_{p \in \mathscr{P}} \sum_{k \in \mathscr{K}_{A}} \lambda_{pk} \left(1 - \left(\frac{\tilde{\lambda}_{a}}{\tilde{\mu}} \right)^{\tilde{Z}_{a}} \right) \prod_{\alpha \in \mathscr{A}} \pi_{\alpha}^{\left(Z_{\alpha} R_{p\alpha a} \right)} \left(\frac{\tilde{\lambda}_{\alpha}}{\tilde{\mu}} \right)^{\left(\tilde{Z}_{\alpha} \beta_{p\alpha a} \right)}$$
(14)

We propose finding the true vehicle utilisations by first solving Equation 13 for the λ_a ; determining the primary vehicle utilisations with Equation 10; then using these to solve Equation 14 for the $\tilde{\lambda}_a$; and determining the secondary vehicle utilisations with Equation 11. These can be solved numerically. In our implementation this is solved using the MINPACK hybrd and hybrj algorithms by using Scipy's fsolve function (Virtanen et al., 2020). These algorithms are efficient methods for finding zeros of systems of non-linear equations. They are based on the Powell hybrid method (Powell, 1970). It should be noted that in this section we assume that all λ_{pk} are static and not time-dependant, which is an assumption that will be used throughout

the optimisation methodology, however in the simulation we can account for time-dependent demand.

Now note that, in Equation 13, we have that

$$\sum_{p \in \mathscr{P}} \sum_{k \in \mathscr{K}} \left(1 - \left(\frac{\lambda_a}{\mu}\right)^{Z_a} \right) \prod_{\alpha \in \mathscr{A}} \left(\frac{\lambda_\alpha}{\mu}\right)^{\left(Z_\alpha \beta_{p\alpha a}\right)} < 1,$$

and so $\sum_{a \in \mathscr{A}} \lambda_a < \sum_{p \in \mathscr{P}} \sum_{k \in \mathscr{K}} \lambda_{pk}$, meaning that there is lost demand. This corresponds to the patients who are abandoned, or take private or public transport to a hospital instead of waiting for an ambulance. Similarly, in Equation 14, we have that

$$\sum_{p \in \mathscr{P}} \sum_{k \in \mathscr{K}_A} \left(1 - \left(\frac{\tilde{\lambda}_a}{\tilde{\mu}} \right)^{\tilde{Z}_a} \right) \prod_{\alpha \in \mathscr{A}} \pi_\alpha^{\left(Z_\alpha R_{p\alpha a} \right)} \left(\frac{\tilde{\lambda}_\alpha}{\tilde{\mu}} \right)^{\left(\tilde{Z}_\alpha \beta_{p\alpha a} \right)} < 1,$$

and so $\sum_{a \in \mathscr{A}} \tilde{\lambda}_a < \sum_{p \in \mathscr{P}} \sum_{k \in \mathscr{K}_A} \lambda_{pk}$. This lost demand represents both abandoned calls, but also calls where secondary vehicles are not deployed as a primary vehicle could reach them first.

4.5. Metaheuristic Optimisation

In order to maximise the MESLMHPHF objective function presented in Section 4.2, we use an evolutionary algorithm. A population of possible solutions is created and ranked according to the objective, then for each generation of the metaheuristic a proportion of the best performing solutions is kept, and are mutated to complete the population for the next generation. To encourage exploration, the number of times each solution is mutated begins high and gradually decreases throughout the run of the simulation. The metaheuristic takes five hyper-parameters: N the population size, κ the number of solutions to retain to the next generation, m_0 the initial mutation rate; c the cooling rate; and H the number of generations. The metaheuristic is described in Algorithm 1.

Mutations consist of one of four possible changes: randomly change the location of one primary vehicle; randomly change the location of one secondary vehicle; randomly remove one primary vehicle and add secondary vehicles to three randomly chosen locations (as one primary vehicle costs the equivalent to three secondary vehicles); and randomly remove three secondary vehicles and add a primary vehicle to a randomly chosen location. If it is required that exact vehicle numbers need to be retained, then only for first two mutations need be chosen.

The source code for the optimisation metaheuristic, including the evaluation of the MESLMHPHF objective function and the solution method for the utilisation considerations, is available at Knight and Palmer (2023) and development took place at https://github.com/drvinceknight/HeterogeneousAmbulanceFleetAllocations. Numerical results are discussed in Section 6.

Algorithm 1 Evolutionary Algorithm used to find better allocations.

- 1: Create a population of *N* random allocations (Z_a, \tilde{Z}_a)
- 2: for $i \leftarrow 0$ to H do
- 3: Rank the population by decreasing expected survival according to Equation 5
- 4: Keep the top κ solutions
- 5: $m \leftarrow [m_0 \min(m_0, i \times c)]$
- 6: **for** $j \leftarrow 0$ **to** $N \kappa$ **do**
- 7: Randomly choose and copy a solution from the population
- 8: Mutate that copy of the solution *m* times
- 9: Place mutated copy in a separate population
- 10: end for
- 11: Combine populations;
- 12: end for
- 13: Rank the population according to Equation 5;
- 14: Output the top ranking solution.

5. Simulating Allocations

Discrete event simulation (DES) is a common methodology that allows us to investigate given scenarios under uncertainty. It is a common methodology for modelling emergency medical services, with a review given in Aboueljinane et al. (2013). Here it will be used to quantify the effectiveness of given ambulance allocations by measuring a range of key performance indicators (KPIs), such as the average response time, the percentage of abandoned calls, vehicle utilisations, and expected survival based on response times. The simulation is built using the Ciw library (Palmer et al., 2019), and the model logic is described below. Its central ideas include modelling transit jobs as customers, rather than the patients themselves, and simulating primary and secondary vehicles as two simulations sequentially, with the output of the former being the input of the latter. Separating out primary and secondary vehicle logic simplifies the model logic, allows for modular simulations that can be run in isolation, and so easier to adapt and to maintain.

5.1. Simulating Primary Vehicles

For primary vehicles the logic to simulate is as follows: Patients make a call from one of a number of pick-up locations and await an ambulance to pick them up and take them to an appropriate hospital. Ambulances are stationed at a number of ambulance locations; when a patient makes a call, all free ambulances from any location calculate their expected time to reach that patient, and the ambulance with the smallest expected time to the patient is called out. The ambulance drives from its current ambulance location to the patient's pick-up location, then if a hospital is required, from the patient's pick-up location to the hospital, and then from

the nearest hospital back to their original ambulance location. If a hospital is not required, the ambulance returns to their original ambulance location.

Rather than considering patients as customers in a queue, here the situation is re-framed to model transit jobs as customers, with the ambulances as servers. This is similar to the approach of Kergosien et al. (2015), whereas ambulances are now simple servers, their stations are nodes in a queueing network, with a routing decision representing the ambulance preference.

Transit jobs can be categorised into classes corresponding to the pick-up locations \mathscr{P} and speciality \mathscr{K} . Jobs of class $(p,k) \in \mathscr{P} \times \mathscr{K}$ arrive with rate λ_{pk} . Servers can be categorised into classes corresponding to the ambulance locations \mathscr{A} . Service times are server-dependent, that is the service time of a transit job of type (p,k), being served by a server of class a, will have service time T_{paky} if a hospital is required and T_{pak} if no hospital is required. These were given by Equations 1 and 2 respectively, in Section 3. That is T_{paky} and T_{pak} is the overall time the ambulance spends dealing with that transit job and is unavailable to receive any more transit jobs.

From a patient's point of view their service only lasts $G_k + C_{py} + J_{k}$, (or G_k if no hospital is required), and their waiting time is B_{pa} plus the time waiting for an ambulance to be dispatched. However in our case, if there are no free ambulances at the time of call, it is assumed that patients find their own care or transport, and so the call is abandoned. Therefore, the time waiting for dispatch is always zero.

Furthermore, all service times are time dependent, and calculated from travel distances and approximate hourly traffic levels, described in Section 5.2.

5.2. Travel Time Calculations

Travel times within a city such as Jakarta are not constant throughout the day due to the variability in traffic, and it is necessary to capture these in models of emergency medical services (Schmid and Doerner, 2010). Here, each day is split into a set of periods, \mathscr{H} indexed by h. Within each period traffic levels are modelled as piecewise linear functions similar to Horn (2000), and are here considered constant but differing from period to period. Traffic levels influence the speed of the ambulance through a delay factor d_h associated with each period. Therefore, the expected speed S_h at which an ambulance travels during period h is given by Equation 15,

$$S_h = sd_h \tag{15}$$

where s is some given baseline speed. The relationship between time travelled and distance covered is piece-wise linear with slope S_h in period h. Therefore, travel times are calculated using this relationship.

As an example, consider the scenario where for the first period $t \in (0,4)$ we have $S_1 = 1$; for $t \in (4,8)$ we have $S_2 = \frac{1}{4}$, for $t \in (8,12)$ we have $S_3 = 2$, and for $t \in (12,16)$ we have $S_4 = 1$. Consider that the vehicle must travel 11 units. If the vehicle begins its journey at t = 0, then the expected travel time would be 11 time units, shown in Figure 4a. However if the



vehicle begins its journey at time t = 3 then the expected travel time is 10 time units, shown in Figure 4b.

This is the method used to calculate the expected values of B_{pa} , C_{py} , D_{ya} and F_{pa} . Furthermore it is assumed that travel times for each segment of the route (from vehicle location to patient, from patient to hospital, and from hospital back the the vehicle's location) follow a Triangular distribution around the calculated expected travel time, between 75% and 125% of the value. A Triangular distribution is a useful distribution to use when not much is known about the travel times except some basic summary statistics (Robinson, 2014), in this case a likely range of values.

5.3. Time-dependent Demand

Additionally, it is assumed that calls arrive according to a Poisson distribution, with rates λ_{pk} , and that these rates are time dependent. Each day is split into four 6-hour periods, the morning, afternoon, evening and night, and each speciality *k* and pick-up location *p* will have a different call arrival rate for each of these periods.

For some locations and specialities the number of observed calls might be very low, and so the λ_{pk} rate would be very low. This can cause synchronicity issues when sampling arrivals (Pidd, 2004), where artificially long inter-arrival times can be introduced at the beginning of one time period due to the low arrival rate in the previous time period. In order to overcome this here, rather than sample inter-arrival times iteratively from an Exponential distribution, an entire schedule of arrival dates are sampled at the beginning of the simulation run, first by sampling the number of arrivals required in each time period from a Poisson distribution. This mechanism, and alternative to thinning (Lewis and Shedler, 1979), was implemented in the Ciw software in version v2.3.3 (The Ciw library developers., 2022), and described in the documentation here: https://ciw.readthedocs.io/en/latest/Guides/time_dependent.html.

5.4. Simulating Secondary Vehicles

Secondary non-transit vehicles, can in general travel faster than primary transport vehicles, and so can be dispatched at the same time as a primary vehicle but reach the patient earlier, reducing the response time for that patient and so increasing their survival probability. They are only dispatched for patients of speciality $k \in \mathscr{K}_A$, and if the closest secondary vehicle can reach them before the closest primary vehicle. Here secondary vehicles are simulated in a second discrete event simulation, run sequentially, but simulating the exact same time period and events as the first simulation. This is similar to sequential hybrid simulation methodology (Brailsford et al., 2019, Morgan et al., 2017), however in this case the two combined components are both DES, and are combined to simplify the logic of each component, maintain modularity and so ease model reusability and future adaptations.

This is possible as there are no synchronicity issues between the two components. Primary vehicles operate independently of secondary vehicles, that is, the way primary vehicles respond to a patient is unaffected by the presence or lack of a secondary vehicle. Therefore, primary vehicle logic is not compromised by simulating primary vehicles in isolation. Secondary vehicles are impacted by the behaviour of primary vehicles, they must remain with the patient until the primary vehicle arrives, and so must be simulated after the simulation of primary vehicles, taking as inputs the exact list of events that occurred. That is, the logic of the secondary vehicles is determined by observing the actions of primary vehicles and reacting to them. Simulation results are combined for each individual patient, to determine their response time caused by either primary or secondary vehicles.

Secondary vehicles are chosen in the same way as primary vehicles, by choosing out of the currently free vehicles the one with the smallest expected time to patient. Their service times are reactionary to what occurred with the primary vehicle. The decision to send a primary vehicle or not depends on the previously estimated travel time for the primary vehicle, not the actual travel time experienced by that primary vehicle: if a secondary vehicle is estimated to reach there first, they will be dispatched. Let \tilde{B}_{pa_2} be the time it takes for the secondary vehicle to travel from its location a_2 to the patient pick-up location p; B_{pa_1} is the time is took for the primary vehicle to get from its location a_1 to the patient pick-up location p; G_k is the time the primary vehicle spent with the patient; and \tilde{F}_{pa_2} is the time it takes to return to the secondary vehicle's location a_2 from the patient pick-up location p. Note that B_{pa_1} and G_k are exact values obtained from the initial simulation of the primary vehicles, while B_{pa_2} and F_{pa_2} are random variables to sample in the subsequent simulation. Travel times are calculated as described in Section 5.2, replacing the primary vehicle delay factor, d_h , with a delay factor for secondary vehicles, d_h . This accounts for secondary vehicles travelling faster and reacting to traffic differently to primary vehicles, similar to the method used in Henderson and Mason (2004) to differentiate the travel times of ambulances with and without flashing lights.

Exact synchronicity considerations are shown in Figure 5:

whenever the secondary vehicle reaches the patient before the primary has left, vehicle, they remain with the patient until the primary vehicle leaves;



FIG. 5. Visualising secondary vehicle (SV) logic reacting to the primary vehicle's (PV) actions.

- if secondary vehicle arrives after the primary vehicle has left, they will immediately return to their stations as they are not needed at the scene;
- if the expected time for the secondary vehicle to reach the patient exceeds the expected time for the primary vehicle to reach the patient then the secondary vehicle is not deployed, and so that transit job would be abandoned in the secondary simulation (although still seen by a primary vehicle in the initial simulation). We assume that secondary vehicles must still reach the pickup locations, due to potential difficulties in communication en-route. Consider a situation when a secondary vehicle arrives and the patient does not require transportation to a hospital, with better en-route communication this would mean the primary vehicle turns around, however without en-route communication which is assumed here, the primary vehicle must still reach the patient pick-up location.

Therefore job service times for secondary vehicles, \tilde{T}_{pak} , are given by:

$$\tilde{T}_{pak} = \max\left(\tilde{B}_{pa_2}, B_{pa_1} + G_k\right) + \tilde{F}_{pa_2}$$
(16)

5.5. Combined Model

Combining the logic presented in Sections 5.1 and 5.4 in a sequential manner, with the output of one determining the input of the other, gives the overall simulation logic for simulating both primary and secondary vehicle types. It is used to quantify the effectiveness of a given allocation (Z_a, \tilde{Z}_a) for all $a \in \mathcal{A}$, by recording several useful KPIs. In this work, the KPIs of interest are: the average primary vehicle utilisation, the average secondary vehicle utilisation, the mean response time, the percentage of abandoned calls (that is a measure of the primary vehicle unavailability), and the expected overall survival. Survival is modelled using a combination of survival function curves and hard cut-offs, and is described in Section 4.1.

Verification of a model is ensuring that the conceptual model is implemented correctly and without error, while validation of a model ensures that the conceptual model accurately represents the real system to be modelled (Robinson, 2014). Verifying the simulation code, alongside verifying the optimisation and MESLMHPHF objective function code, is done through automated code testing (Palmer and Tian, 2023, Percival, 2014). This includes unit testing to ensure that each sub-component works are requires, and is bug and error free; end to end testing, ensuring that the components interact together as intended; and extreme value



FIG. 6. Map of Jakarta's Municipalities.

testing ensuring that changes in input parameters result in sensible and expected changes in output results and behaviours. The testing suites are included in the open code repositories provided.

The source code for the combined simulation, is available at Palmer and Tuson (2023) and development took place at https://github.com/MarkTuson/ambulance_simulation.

6. The Case of Jakarta

Jakarta, the capital of Indonesia, has a population of 10.5 million (Badan Pusat Statistik Provinsi DKI Jakarta, 2023) residing within 664 km². The city is divided into five municipalities and one district, each of which is divided into sub districts and, in turn, neighbourhoods. There are in total 42 sub districts and 261 neighbourhoods within mainland Jakarta (excluding the Thousand Islands regency in the north) (Badan Pusat Statistik Provinsi DKI Jakarta, 2020); a map is given in Figure 6.

As described in Section 1, there are many challenges to calling for an ambulance in Jakarta, as captured in Brice et al. (2022). Based on this work the regional government of Jakarta invested in a new fleet of coordinated ambulances accessible via the single emergency number 119. This work, in collaboration with Ambulans 118 and the 119 Emergency Ambulance Service, finds potential fleet allocations using the optimisations methods given in Section 4, and evaluates these allocations' effectiveness using the simulation described in Section 5.

As of October 2022, the 119 service in Jakarta ran a fleet of 81 primary Emergency Ambulances (EAs) and 13 secondary motorbike Rapid Response Vehicles (RRVs), distributed across 67 ambulance stations throughout the city. Figure 7 summarise the allocations vehicle locations. Working with our ambulance partners, after careful consideration the following patient

0 ON THE OPTIMISATION OF HETEROGENEOUS AMBULANCE FLEET ALLOCATIONS



FIG. 7. Map of current Emergency Ambulance (EA) and Rapid Response Vehicle (RRV) locations and allocations to ambulance bases across Jakarta.

'specialities' ($k \in \mathcal{K}$) were agreed based on clinical need (although other categories could readily be adopted and included as necessary):

- A1 High priority emergency patients: critical patients who require immediate lifesaving assistance with a target ambulance response time of 8 minutes. In our data set there were 168 identified calls that met this criterion. As in Knight et al. (2012), equation (3) is here used to measure survival of these patients.
- A2 Other emergency patients: urgent patients that require assistance with a target response time of 15 minutes. In our data set there were 23,784 calls of this type.
- **B** Non-emergency patients: patients that still benefit from an ambulance response but non-critical with a target response time of 60 minutes. There were 31,725 calls of this type in the data set.

Secondary vehicles can be called to patients of speciality A1 and A2, therefore $\mathcal{H}_A = \{A1, A2\}$ and $\mathcal{H}_B = \{B\}$. The number of calls from each of these specialities varies by neighbourhood, as shown in Figure 8. It can be seen that many calls are highly concentrated in a handful of neighbourhoods rather than spread throughout the city. In this work, we approximate pick-up locations by the geographic centroid of each of the 261 neighbourhoods, each representing an ambulance pick-up that that occurred within that neighbourhood.

Firstly, in Section 6.1 we use the MESLMHPHF objective function of expected survival, and the simulation, to find KPIs that to compare the performance of two proposed grid allocations to the currently used allocation. An additional KPI is also calculated, the expected survival of A1 patients, found by taking Equation 5 and summing over k = A1 only. Then in Section 6.2 we outline four possible future demand scenarios, and in Section 6.3 we use the optimisation



FIG. 8. Number of calls received by speciality and neighbourhood.

model to find and compare improved allocations for the current number of vehicles. Finally in Section 6.5 we investigate the number of vehicles required to meet the demand across the given scenarios.

All models are parameterised using demand data that covers all 261 neighbourhoods from 1 January to 31 December 2019, before the COVID-19 pandemic. Data from 2020-2021 was naturally heavily skewed by the pandemic, with significantly lower demand, and so was not considered as representative of a typical period for forecasting future needs, hence the decision to use the available 2019 demand. APPENDIX C provides further details on the parameterisation of the model, including how travel time estimates were obtained. Verification of the models were discussed in Section 5.5, we validated based on expert opinion from our collaborators, Ambulans 118 and the 119 Emergency Ambulance Service in Jakarta, where several workshops have been carried out in Indonesia where the model was shared and discussed. This expertise was vital in giving a basic validation of the model, with collaborators approving model results parametrised from their provided data and expert opinion.

6.1. Evaluating 'Grid' Proposals

Discussions with senior staff at 119 identified that they were seriously considering a reconfiguration their ambulance allocations into a grid structure, placing an ambulance station at regular intervals throughout the city and uniformly distributing the vehicles that they felt would ensure equitable coverage. Two possible grid allocations were being considered: one placing an EA every 3km across the city (giving a total of 70 vehicles), and another placing three EAs every 5km across the city (giving a total of 72 vehicles). Figure 9 show these proposed allocations.

Running the simulation and MESLMHPHF objective function of expected survival for both the current and the two proposed grid allocations gives the results shown in Table 1. The simulation was run over half a year, with a month warm up time, over five replications. The variability over the replications is very small, and we have included 90% conficence intervals around the mean response times to demonstrate this. We observe that the two proposed grid allocations perform much worse than the current allocation as measured by mean response

22 ON THE OPTIMISATION OF HETEROGENEOUS AMBULANCE FLEET ALLOCATIONS



TABLE 1 Calculated KPIs for the current and proposed grid allocations.

time, survival, and percentage of calls abandoned. This is expected as vehicle numbers are lower under the new proposals, as well as the allocations focussing on coverage rather than response times or survival.

/wnloaded from https://academic.oup.com/imaman/advance-article/doi/10.1093/imaman/dpaf027/8190187 by guest on 22 July 2025

Reason	% of emergency respondents
Used Ambulance	13.0
Too expensive	5.2
Not available	13.1
Would take too long	16.5
Not necessary	15.7
Not aware of service	33.1
Other	3.5

TABLE 2Barriers to use of an ambulance service inJakarta, from Brice et al. (2022).

6.2. Demand Scenarios

Similarly to many other cities worldwide, ambulance services in Jakarta anticipate that demand for services will grow in future years, not least because there is now a coordinated service accessible via a single common number '119' to call. This should help raise awareness and visibility of an ambulance service. In the case of Jakarta, certainly increased use of 119 and therefore increased demand would be a desired outcome and is pro-actively supported by the Indonesian Government.

In other cities with more detailed data collection, statistical models can be used to predict future spacial demand across a city, as in Nicoletta et al. (2022). In the case of Jakarta, collected call data is not detailed enough, but we do consider four different demand scenarios, developed by considering the result of our cross-sectional study of patients attending EDs in Jakarta (Brice et al., 2022). As part of the survey, randomly selected patients arriving at each emergency department were asked whether they had used an ambulance to attend, and reasons for not using an ambulance. For those patients who were categorised by the medical staff as emergency and for whom therefore an ambulance might have seemed a sensible option, the responses are given in Table 2.

The survey confirmed three major barriers to use, namely: the service's visibility (33.1% of respondents were unaware of the service), the service's reliability (13.1% of respondents reported that there was no ambulance available and 16.5% of respondents believed it would take too long), and the service's cost (5.2% of respondents cited cost as a deterrent). Therefore, four demand scenarios were considered, corresponding to addressing each of these barriers in turn:

D13: representing the current situation where approximately 13% of emergency patients (specialities A1 and A2) do use an ambulance.

D19: representing the situation where visibility is addressed. Here, we distribute the 33.1% of the respondents who were unaware of the service proportionally between using the

Parameter	Value Used	Explanation	
μ	$\frac{1}{3.886 \times 60}$	Average primary vehicle service rate (hr^{-1}) .	
$ ilde{\mu}$	$\frac{1}{1.038 \times 60}$	Average secondary vehicle service rate (hr^{-1}) .	
N	100	Population size.	
κ	20	Number of solutions to keep per generation.	\sim
m_0	6	Initial number of mutations.	
С	0.1	Cooling rate, rate at which number of mutations decreases.	
Н	200	Number of generations.	

TABLE 3 Hyper-parameters used in the optimisation for all experiments.

ambulance and amongst the remaining issues. Using this methodology we would expect 19.4% of emergency patients to now use an ambulance.

- D34: representing the situation where reliability and visibility are addressed. Using the same methodology we would expect 34.8% of emergency patients to now use an ambulance.
- **D45**: representing the situation where cost, reliability and visibility are all addressed. Using the same methodology we would expect 45.4% of emergency patients to now use an ambulance.

Guided by our findings in Brice et al. (2022), recent changes mean that the 119 services is now free to use for Jakarta residents, so certainly the highest demand (**D45**) scenario is plausible in the near future. After discussions with our partners, it was agreed that for each of the scenarios we should assume that non-emergency demand (speciality B) remains unchanged. Therefore, in each scenario, we multiply the call demand for each municipality, λ_{pk} by $\frac{13\%}{13\%} = 1$, $\frac{19\%}{13\%} = 1.46$, $\frac{34\%}{13\%} = 2.62$, and $\frac{45\%}{13\%} = 3.46$ respectively, for all $k \in \mathcal{K}_A$, $p \in \mathcal{P}$. Note that for some pick-up locations p, λ_{pk} is very small, possibly zero, which is realistic in scenario **D13**, as that is what is observed, but may result in artificially low demand in those locations when using the multiplication method to scale up demand.

6.3. Improving the Current Allocation

Applying the optimisation metaheuristic from Section 4 to the current allocation re-allocates the 81 EAs and 13 RRVs. We do this for each of the four scenarios described in Section 6.2. The hyper-parameters used for this optimisation algorithm throughout this study are given in Table 3, which include approximations for the average service times of primary and secondary vehicles derived from the initial simulation model. APPENDIX D shows some initial explorations on the hyper-parameter choices, giving confidence that the chosen parameters, with a large enough number of iterations N, are sufficient to find allocations that perform well.

Demand	Allocation	MESLMHPHF Objective function	MESLMHPHF A1 Objective function	Observed Survival	Observed A1 Survival	Primary Utilisation	Secondary Utilisation	Mean Response Time (mins)	Response Time StDev (mins)	Abandoned	
D13	Current	98.34%	26.05%	77.83% (± 0.25%)	17.80% (± 1.68%)	30.99%	22.47%	17.69 (± 0.07)	6.98	0.00%	
	Improved	99.10%	30.99%	77.16% (± 0.28%)	16.22% (± 0.42%)	31.06%	18.29%	17.86 (± 0.05)	7.04	0.00%	
D19	Current	97.74%	25.17%	70.01% (± 0.50%)	15.42% (± 0.65%)	38.11%	32.24%	18.48 (± 0.03)	7.73	0.44%	
	Improved	99.64%	29.85%	69.58% (± 0.56%)	14.35% (± 0.69%)	38.12%	27.39%	18.52 (± 0.05)	7.70	0.46%	4
D34	Current	95.78%	22.75%	54.37% (± 0.14%)	12.53% (± 0.31%)	54.07%	52.28%	22.21 (± 0.03)	11.56	3.44%	
	Improved	99.61%	29.59%	54.84% (± 0.31%)	12.53% (± 0.56%)	54.16%	49.07%	22.13 (± 0.05)	11.86	3.79%	
D45	Current	92.88%	20.67%	47.22% (± 0.21%)	11.86% (± 0.31%)	61.79%	60.89%	23.99 (± 0.04)	12.92	9.69%	
	Improved	99.54%	29.36%	48.32% (± 0.20%)	11.66% (± 0.28%)	61.62%	56.80%	23.18 (± 0.07)	12.51	9.27%	

 TABLE 4
 Calculated KPIs for the current and improved allocations under the four possible demand scenarios.

Table 4 compares the MESLMHPHF objective function value and simulation derived KPIs between the current and the optimised allocations, for each demand scenario, along with 90% confidence intervals over the trials. It can be seen that, as expected, as demand increases then utilisation of primary and secondary vehicles increase, as the ambulance service is busier. Similarly, mean response time increases with demand, this is due to the vehicles being busier, and therefore less chance that the most optimally placed vehicle is dispatched. This is confirmed by the percentage of calls abandoned increasing with demand, showing ambulance unavailability.

The value of the MESLMHPHF objective function and the observed survival, as given by the simulation model, do not match. Section 6.4 discusses this in more detail.

As desired, the optimisation heuristic has indeed improved the allocation as measured by the MESLMHPHF objective function. In particular A1 patients are benefiting from the improved allocations as measured by the MESLMHPHF objective function, though not by the simulation. Why the simulation might not capture this is hypothesised in Section 6.4. This measured improvement is due to these patients being more sensitive to small fluctuations in response time, due to their survival function Equation 3. Also, using a better allocation does increase the survival, and that the drop in survival due to demand increases is not as severe when using an improved allocation in comparison to the current allocation, although more numerical evidence would be helpful here. This improved robustness not only validates the heuristic but also highlights the research contribution of Section 4.3 as an approach to handle the recursive relationship between an allocation x and the utilisation rates.

It is worth noting that for the simulation based KPIs (mean response times and percent abandoned), there is not much difference between the currently used allocation and the better allocation, showing that these KPIs may not be good approximations of survival.

6.4. Comparing Expected & Simulated Survival Measures

In Table 4 it can be seen that there is a large discrepancy between the expected survival, both overall and of A1 patients, computed using the MESLMHPHF described in Section 4.2 and calculated using the simulation. This is a consequence of the variability in response times, which is explicitly modelled by the simulation but not captured by the MESLMHPHF. As seen in the table, the standard deviations of the response times are relatively large in comparison to the mean response times, and so this effect would be substantial here.

The discrepancies are exacerbated by the use of cut-off survival function curves for most of the patients: consider a survival function with cut-off L, with response time below L the patient's chance of survival is 0, and above L the patient's chance of survival is guaranteed. The optimisation will find a solution aiming to push the expected response time above L, and once this is reached, possibly with an expected response time as low as L itself, it will not attempt to push the expected response time above this, as it will have no effect on the expected survival. However considering variability of mean response times, a considerable proportion of the response times will lie below the expected value, and so possibly below L, resulting in worse survival probabilities, as observed in Table 4. Therefore the distributions chosen to model travel times can have a large impact on the survival measures, in this case the naive Triangular distribution.

Therefore we present two separate models to measure overall survival, both of which have drawbacks in terms of capturing the stochasticity of the system. We present both as a holistic view of the system, and suggest the true outcomes lie somewhere in between the two measures. As the heuristic improves the outcomes in one measure and does no harm in the other, we anticipate that this procedure will improve patient outcomes. This leads to a number of avenues of future work to improve this and overcome this effect. Variability could be explicitly modelled within the MESLMHPHF, such that it is not only maximising *expected* survival, but considers some percentiles of guaranteed survival. Another avenue would be to use the simulation explicitly in the optimisation process, incorporating the simulation into the objective function and using simulation-optimisation techniques to improve runtime efficiency.

Due to the naive choice of travel time distributions, including difficulties in being able to validate such a model and the choices made, we consider only the MESLMHPHF objective from now on.

6.5. Finding Better Allocations for Any Resource Level

The metaheuristic algorithm described in Section 4.5 finds allocation of ambulances across the 67 current ambulance stations for a given number of primary and secondary vehicles. Information supplied by one of the ambulance operators suggested that in terms of total running costs one EA (primary vehicle) was approximately equivalent to three RRVs (secondary vehicles), and so we consider three RRVs to be one resource; that is, a resource level of 60 could represent 60 primary and 0 secondary vehicles, or 59 primary and 3 secondary vehicles, or 58

primary and 6 secondary vehicles, and so on. For a given resource level we run the optimisation algorithm allowing for one primary to be swapped for three secondary vehicles, and vice versa, and we report KPIs for the best performing, in terms of maximal survival, combinations, as well as for the case when the number of secondary vehicles is fixed at zero. We call these the multiple vehicle type and single vehicle type scenarios.

Figures 10a-15b display the obtained KPIs for each of these allocations. Confidence intervals over the replications were too small to display on the plots. It can immediately be seen that increasing the resource level has a positive effect on all KPIs, with vehicle utilisations, mean response times, and percentage abandoned decreasing, and overall survival increasing. Similarly, as expected increasing the demand of emergency calls has a negative impact on all KPIs.



y-axis scale in order to keep the data points visible.)

It is particularly interesting to compare the scenario in which single vehicle types (only emergency ambulances) were allocated against the scenario where multiple vehicle types (both emergency ambulances and rapid response vehicles) were allocated. For equivalent resource levels, introducing secondary vehicles increases ambulance utilisation, and so decreases availability, resulting in an increase in abandoned calls. However, introducing secondary vehicles also results in a decrease in the mean response time, and a large increase in the expected survival, especially those of speciality A1. In fact, it can be seen from Figure 11b that when using multiple vehicle types the survival of A1 patients plateaus as the resource level increases, indicating that the system has reached an upper bound on the survival of A1 patients. This is corroborated by considering the survival function of A1 patients given in Equation 3, where even if all patients were seen immediately a survival of 43.5% would be expected, and



so accounting for even a minimum travel time to the patients an observed plateau value of 32% could not be improved upon by adding more resource level.

An interesting phenomenon occurs when looking at Figure 14b. Here it seems that, when using multiple vehicles, mean response time decreases as demand increases, which is counterintuitive. This is however due to a peculiar interplay between demand and vehicle behaviours: as stated in Section 6.2 only demand for emergency patients are increased (specialities A1 and A2), which are also the patients that can be seen by secondary vehicles, and so would have a lower response time than non-emergency patients (speciality B). Thus, demand scenario D45 receives a larger proportion emergency patients than scenario D13, and so a larger proportion of patients with lower response times due to being seen by secondary vehicles; bringing the



overall mean response time down. This does not mean however that emergency patients in scenario D45 have lower response times than those in scenario D13.

It is also noticeable that the increases or decreases in the simulation derived KPIs are not necessarily monotonic with increases in resource level. This is due to the derived allocations being based on the MESLMHPHF score only, which is a measure of survival, while the simulation derived KPIs are indirect measures of performance. Again, this may indicate that these indirect measures of performance, such as mean response times or percentage within target, are not good indicators of survival.

Crucially, the plots show that the allocations produced by the optimisation algorithm perform better than the current allocation for the same demand and resource levels. The current levels include 81 EAs and 13 RRVs, so a resource level of 85.33, with multiple vehicles. Comparing the derived allocations for this level in Figures 10a-15b to the results given in Table 1,



FIG. 15. Percent of abandoned calls results

the derived allocations give lower ambulance and RRV utilisations, lower mean response times, and a higher percentage of patients seen within target.

7. Discussion & Conclusions

This paper has described the development of EMS demand and capacity models and their application to the city of Jakarta, Indonesia. To our knowledge, this is the first such study of analysing emergency ambulance services in Jakarta, and the work has directly assisted and guided providers and the government to make investment decisions for a coordinated and free to use EMS.

Firstly, we have considered patient survival and outcomes within a developed MESLMH-PHF model. By considering geospatial travel times and accounting for the varying needs of different patient types through categorising patients by speciality and through using different survival function profiles, an expected overall survival was given. This was used within a optimisation processes, with an evolutionary metaheuristic, to find vehicle fleet allocations that maximised expected survival. This expected survival function extends previous work in Knight et al. (2012) to include more than one vehicle type. A particular novelty here was numerically solving implicit utilisation relationships, equations 13 and 14, to calculate the expected survival, as an alternative to busy fractions and hypercube models.

Secondly, a discrete event simulation model has been used to evaluate existing and potential heterogeneous vehicle ambulance feet allocations in terms of key performance measures, such as response times, survival, and vehicle utilisation rates. This takes into consideration geospatial demand and travel, and temporal variation in demand and traffic levels. A novel feature of our approach is that the model comprises of sequential simulations, feeding data directly from one into the other in order to simulate primary and secondary vehicles separately while maintaining synchronicity, with the overall aim of reducing the complexity of the simulation logic and maintaining modularity.

Both models investigate ambulance service activity in the case of heterogeneous patients and heterogeneous fleets. Heterogeneous patient groups consider those with distinct demand profiles, priorities, and survival functions. Heterogeneous fleets concern different types of vehicles, in our case emergency ambulances (EAs) which respond to every patient, and Rapid Response Vehicles (RRVs) which can be utilised to reach patients faster, despite being unable to transport patients themselves.

Using a combination of models has permitted an approach that can capture performance measures or take into account factors that each model in isolation can not. For example, the maximum expected survival model is a lot more appropriate for use within an optimisation process as the runtimes are a far more reasonable than the simulation. On the other hand, the simulation allows for greater complexity such as temporal demand, and is able to capture a greater range of KPIs. Crucially, using mixed methodology such as this allowed for independent evaluation of the output of one using the other; but also insights from an initial run of the simulation gave helped parametrise and justify assumptions of the optimisation.

Combinations of models allow for robust conclusions to be drawn about the system, by considering the specific attributes of each modelling approach. Here, the MESLMHPHF is quick to run, but can fail to capture stochasticity, while the simulation is parametrised using expert opinions that can effect its robustness. Taking both into account offers a better approximation and more insight.

A key feature of both models, and their novelty, is the consideration of heterogeneous fleets. In a highly populated area such as Jakarta (around 16,000 population per km²), the inclusion of RRVs such as paramedics on motorbikes is very important, given RRVs can access areas that cannot easily or quickly be accessed by ambulances. Results from both the optimisation and simulation research showed that the RRVs can help increase the overall survival and reduce the response times, especially crucial when responding to life threatening events such as cardiac arrests (Holmén et al., 2020).

In building the models, some assumptions or relaxations were made to simplify the model, and future work should investigate the effects of these, for example the assumption that service times are not location dependent is unrealistic, but relaxed here. Further developments of the model could consider explicitly modelling service rates using the travel time matrices. Another limitation identified was the lack of consideration of variability in travel times in the optimisation model, which cased some discrepancies between the simulated and expected survivals. One method to overcome this could be to incorporate variability directly into the MESLMHPHF objective, for example using percentiles instead of averages, or to optimise using the simulation as an objective. However in both approaches, a good estimate of travel time distributions would be needed, with better validation. Another direction of future work is to find exact methods of solving the MESLMHPHF instead of relying on a metaheuristic algorithm.

The results of our research also suggest that allocation strategies may not be intuitive to ambulance service managers, further emphasising the value of a modelling approach. For example, senior managers suggested the use of a grid allocation to maximise geographic coverage, with ambulances equally spread across the city. This was shown to be sub-optimal because neighbourhoods have different population densities and characteristics. For example, the number of daily commuters in Jakarta is considerably high during day time due to work and education (Badan Pusat Statistik Provinsi DKI Jakarta, 2001). Municipalities where trading, educational institutions, and offices are concentrated may become more populated during day time. The model has quantified the deterioration in response times and patient outcomes should a fixed grid system be implemented, leading to 119 to drop this consideration.

The ambulance posts in Jakarta depend on the service provider. Those ambulances provided by the government are in various locations including government buildings as well as community clinics and sub-district hospitals. In our case, we used current ambulance posts as the locations for allocating the resources. As demands for ambulances may change from time to time, future work could evaluate different, and perhaps dynamic, ambulance posts that depend on changing demand volumes in neighbourhoods by time of the day.

The data used for ambulance demand covered one year from 1 January to 31 December 2019, prior the COVID-19 pandemic. Recent studies related to ambulance demand indicate that the pandemic has severely impacted on the utilisation of EMS, for examples in call volumes (San et al., 2021) and in specific medical conditions such as trauma (Azbel et al., 2021), and possibly still continue to affect demand. Ambulance providers, such as 118 and 119, may use our modelling tools to incorporate future demand data and re-evaluate resource needs. To aid this, we are currently working on interfacing the simulation and optimisation into a single easy to use decision support tool with a dashboard (data visualiser).

Studies have shown that the quality in pre-hospital data collection varies considerably in Indonesia (Hooper et al., 2019). We encountered similar challenges and have suggested to 119 and the Indonesian Government that they should continue to strive to improve the coverage and quality of data collection. Future work could explore in more depth unmet demand that is not recorded in the ED or in the ambulance data.

Indonesia is a vast country with an uneven population density and varying quality of available healthcare resources. The models have been parameterised with data from a high population, high density urban area which is also relatively well resourced with healthcare facilities compared to other regions in the country. Ambulance services in rural areas of Indonesia may not have the same quality in management and organisation compared to the capital, Jakarta. We intend to conduct future studies to apply the developed models to analyse ambulance demand and allocations in different areas of Indonesia, including in rural regions.

In conclusion, the developed models have demonstrated a novel approach in modelling ambulance allocations that incorporate multiple vehicle types and health conditions, whilst





B. Full MESLMHPHF Optimisation Model

Decision variables: Z_a , \tilde{Z}_a , Maximise:

$$g\left(Z_{a}, ilde{Z}_{a}
ight) = \sum_{p \in \mathscr{P}} \sum_{a \in \mathscr{A}} \left(\sum_{k \in \mathscr{K}_{A}} w_{k} \lambda_{pk} \hat{\Psi}_{kpa} + \sum_{k \in \mathscr{K}_{B}} w_{k} \lambda_{pk} \Psi_{kpa}
ight)$$

where

$$\Psi_{kpa} = s_k \left(b_{pa} \right) \left(1 - \pi_a^{Z_a} \right) \prod_{\alpha \in \mathscr{A}} \pi_\alpha^{\left(Z_\alpha \beta_{p\alpha a} \right)}$$

and

$$\begin{split} \hat{\Psi}_{kpa} &= s_k \left(\tilde{b}_{pa} \right) \left(1 - \tilde{\pi}_a^{\tilde{Z}_a} \right) \prod_{\alpha \in \mathscr{A}} \tilde{\pi}_{\alpha}^{\left(\tilde{Z}_{\alpha} \beta_{p\alpha a} \right)} \pi_{\alpha}^{\left(Z_{\alpha} R_{p\alpha a} \right)} \\ &+ s_k \left(b_{pa} \right) \left(1 - \pi_a^{Z_a} \right) \prod_{\alpha \in \mathscr{A}} \pi_{\alpha}^{\left(Z_{\alpha} \beta_{p\alpha a} \right)} \tilde{\pi}_{\alpha}^{\left(\tilde{Z}_{\alpha} \left(1 - R_{pa\alpha} \right) \right)} \end{split}$$

 $Z_a, \tilde{Z}_a \in \mathbb{N}$

subject to

where \mathscr{L} is the resource level.

C. Model Parameters

Call arrival rates λ_{pk} are derived from 2019 demand data split by municipality and speciality, and time of day. Probabilities q_{pky} are similarly derived from the 2019 data of transit journeys. All traffic-free travel times are found using Google Maps API, while time-dependent traffic delays are found from the TomTom website and given in Table C.5.

The other models are parameterised in the following way:

• From a sample of ealls the time at site G_k was found to follow a lognormal distribution with parameters $\mu = -0.6219$ and $\sigma = 0.8048$. For the case of Jakarta this was modelled

h	0000-0500	0500-1500	1500-1800	1800-0000
d_h	0.98	0.66	0.59	0.77
\tilde{d}_h	1.96	0.91	0.83	1.16

TABLE C.5 Primary and secondary delay factors.

36 ON THE OPTIMISATION OF HETEROGENEOUS AMBULANCE FLEET ALLOCATIONS

identically for all specialities k. Comparison between the lognormal fit and the sampled times are given in Figure C.1.

- From discussions with staff at the ambulance service in Jakarta, the time at hospital J_k was modelled as a Uniform distribution between 40 and 60 minutes for the emergency specialities A1 and A2, and between 20 and 30 minutes for non-emergency speciality B.
- From discussions with staff at the ambulance service in Jakarta, the refill time Θ is taken to be 60 minutes for an emergency ambulance, and 15 minutes for an RRV.

D. Exploration of Optimisation Hyperparameters

Figure D.2 shows the performance of the evolutionary metaheuristic algorithm under different values of the hyperparameters N, κ , m_0 , and c.

E. Improved Allocations for Current Vehicle Numbers

Figures E.3, E.4, E.5 and E.6 show the improved allocations for 81 primary and 13 secondary vehicles under demand scenarios **D13**, **D19**, **D34**, and **D45** respectively.



FIG. C.1. Comparison between the sampled time on site and the lognormal fit.





FIG. E.3. Improved allocation of 81 EAs and 13 RRVs, under scenario D13.

38 ON THE OPTIMISATION OF HETEROGENEOUS AMBULANCE FLEET ALLOCATIONS



FIG. E.4. Improved allocation of 81 EAs and 13 RRVs, under scenario D19.



FIG. E.5. Improved allocation of 81 EAs and 13 RRVs, under scenario D34.



FIG. E.6. Improved allocation of 81 EAs and 13 RRVs, under scenario D45.

REFERENCES

- Aboueljinane, L., Sahin, E., and Jemai, Z. (2013). A review on simulation models applied to emergency medical service operations. *Computers & Industrial Engineering*, 66(4):734–750.
- Amorim, M., Antunes, F., Ferreira, s., and Couto, A. (2019). An integrated approach for strategic and tactical decisions for the emergency medical service: Exploring optimization and metamodel-based simulation for vehicle location. *Computers & Industrial Engineering*, 137:106057.
- Aringhieri, R., Bruni, M., Khodaparasti, S., and van Essen, J. (2017). Emergency medical services and beyond: Addressing new challenges through a wide literature review. *Computers & Operations Research*, 78:349–368.
- Azbel, M., Heinänen, M., Lääperi, M., and Kuisma, M. (2021). Effects of the covid-19 pandemic on trauma-related emergency medical service calls: a retrospective cohort study. *BMC emergency medicine*, 21(1):1–10.
- Badan Pusat Statistik Provinsi DKI Jakarta (2001). *Migrasi Penduduk JABOTABEK 2001*. Badan Pusat Statistik provinsi DKI Jakarta.
- Badan Pusat Statistik Provinsi DKI Jakarta (2020). *DKI Jakarta Province in Figures*. Badan Pusat Statistik provinsi DKI Jakarta.
- Badan Pusat Statistik Provinsi DKI Jakarta (2023). Analisis Profil Penduduk Provinsi DKI Jakarta: mendeskripsikan Peran Penduduk dalam Pembangunan. Badan Pusat Statistik provinsi DKI Jakarta.
- Bélanger, V., Lanzarone, E., Nicoletta, V., Ruiz, A., and Soriano, P. (2020). A recursive simulationoptimization framework for the ambulance location and dispatching problem. *European Journal of Operational Research*, 286(2):713–725.
- Bélanger, V., Ruiz, A., and Soriano, P. (2019). Recent optimization models and trends in location, relocation, and dispatching of emergency medical vehicles. *European Journal of Operational Research*, 272(1):1–23.
- Birmingham, L., Arens, A., Longinaker, N., and Kummet, C. (2021). Trends in ambulance transports and costs among medicare beneficiaries, 2007–2018. *The American Journal of Emergency Medicine*, 47:205–212.
- Boutilier, J. and Chan, T. (2022). Drone network design for cardiac arrest response. *Manufacturing & Service Operations Management*, 24(5):2407–2424.
- Brailsford, S., Eldabi, T., Kune, M., Mustafee, N., and Osorio, A. (2019). Hybrid simulation modelling in operational research: A state-of-the-art review. *European Journal of Operational Research*, 278(3):721–737.
- Brice, S., Boutilier, J., Gartner, D., Harper, P., Knight, V., Lloyd, J., Pusponegoro, A., Rini, A., Turnbull-Ross, J., and Tuson, M. (2022). Emergency services utilization in Jakarta (Indonesia): a cross-sectional study of patients attending hospital emergency departments. *BMC health services research*, 22(1):639-639.
- Chang, C., Abujaber, S., Reynolds, T., Camargo, C. J., and Z. O. (2016). Burden of emergency conditions and emergency care utilization: New estimates from 40 countries. *Emergency Medical Journal*, 33:794–800.
- Daskin, M. (1983). A maximum expected covering location model: formulation, properties and heuristic solution. *Transportation science*, 17(1):48–70.
- Erkut, E., Ingolfsson, A., and Erdogan, G. (2008). Ambulance location for maximum survival. *Naval Research Logistics*, 55(1):42–58.

- Farahani, R., Fallah, S., Ruiz, R., Hosseini, S., and Asgari, N. (2019). Or models in urban service facility location: A critical review of applications and future developments. European journal of operational research, 276(1):1–27.
- Fraser, A., Newberry Le Vay, J., Byass, P., et al. (2020). Time-critical conditions: assessment of burden and access to care using verbal autopsy in agincourt, south africa. BMJ Global Health, 5:e002289.
- Henderson, S. and Mason, A. (2004). Ambulance service planning: simulation and data visualisation. Operations research and health care: a handbook of methods and applications, pages 77–102.
- Holmén, J., Herlitz, J., Ricksten, S., Strömsöe, A., Hagberg, E., Axelsson, C., and Rawshani, A. (2020). Shortening ambulance response time increases survival in out-of-hospital cardiac arrest. Journal of the American Heart Association, 9(21):e017048.
- Hooper, C., Ranse, J., and Hutton, A. (2019). How is ambulance patient care and response time data collected and reported in malaysia and indonesia? Australasian Journal of Paramedicine, 16:1-8.
- Horn, M. (2000). Efficient modeling of travel in networks with time-varying link speeds. Networks: An International Journal, 36(2):80-90.
- Kergosien, Y., Bélanger, V., Soriano, P., Gendreau, M., and Ruiz, A. (2015). A generic and flexible simulation-based analysis tool for ems management. International Journal of Production Research, 53(24):7299-7316.
- Knight, V., Harper, P., and Smith, L. (2012). Ambulance allocation for maximal survival with heterogeneous outcome measures. Omega, 40(6):918–926.
- Knight, V. and Palmer, G. (2023). drvinceknight/heterogeneousambulancefleetallocations: v1.0.0.
- Lewis, P. and Shedler, G. (1979). Simulation of nonhomogeneous poisson processes by thinning. Naval Research Logistics Quarterly, 26(3):403–413.
- Li, X., Zhao, Z., Zhu, X., and Wyatt, T. (2011). Covering models and optimization techniques for emergency response facility location and planning; a review. Mathematical Methods of Operations Research, 74:281-310.
- Liu, Y., Yuan, Y., Shen, J., and Gao, W. (2021). Emergency response facility location in transportation networks: A literature review. Journal of traffic and transportation engineering (English edition), 8(2):153-169.
- Lowthian, J., Cameron, P., Stoelwinder, J., Curtis, A., Currell, A., Cooke, M., and McNeil, J. (2011). Increasing utilisation of emergency ambulances. Australian Health Review, 35(1):63-69.
- McCormack, R. and Coates, G. (2015). A simulation model to enable the optimization of ambulance fleet allocation and base station location for increased patient survival. European Journal of Operational Research, 247(1):294-309.
- McLay, L. A. and Mayorga, M. E. (2010). Evaluating emergency medical service performance measures. Health care management science, 13:124–136.
- Morgan, J., Howick, S., and Belton, V. (2017). A toolkit of designs for mixing discrete event simulation and system dynamics. European Journal of Operational Research, 257(3):907–918.
- Mukhopadhyay, A., Pettet, G., Vazirizade, S., Lu, D., Jaimes, A., El Said, S., Baroud, H., Vorobeychik, Y., Kochenderfer, M., and Dubey, A. (2022). A review of incident prediction, resource allocation, and dispatch models for emergency management. Accident Analysis & Prevention, 165:106501.
- Nicoletta, V., Guglielmi, A., Ruiz, A., Bélanger, V., and Lanzarone, E. (2022). Bayesian spatio-temporal modelling and prediction of areal demands for ambulance services. IMA Journal of Management Mathematics, 33(1):101-121.

- Palmer, G., Knight, V., Harper, P., and Hawa, A. (2019). Ciw: An open-source discrete event simulation library. *Journal of Simulation*, 13(1):68–82.
- Palmer, G. and Tuson, M. (2023). Marktuson/ambulance_simulation: v1.0.0.
- Palmer, G. I. and Tian, Y. (2023). Implementing hybrid simulations that integrate des+ sd in python. *Journal of Simulation*, 17(3):240–256.
- Percival, H. (2014). Test-driven development with Python. "O'Reilly Media, Inc.".
- Pidd, M. (2004). Computer simulation in management science. Number 5th. John Wiley and Sons Ltd.
- Powell, M. J. (1970). A hybrid method for nonlinear equations. Numerical methods for nonlinear algebraic equations, pages 87–161.
- Pusponegoro, A. (2003). Terrorism in indonesia. Prehospital and disaster medicine, 18(2):100-105.
- Reuter-Oppermann, Mand van den Berg, P. and Vile, J. (2017). Logistics for emergency medical service systems. *Health Systems*, 6(3):187–208.
- Robinson, S. (2014). Simulation: the practice of model development and use. Bloomsbury Publishing.
- Şan, İ., Usul, E., Bekgöz, B., and Korkut, S. (2021). Effects of covid-19 pandemic on emergency medical services. *International Journal of Clinical Practice*, 75(5):e13885.
- Schmid, V. and Doerner, K. (2010). Ambulance location and relocation problems with time-dependent travel times. *European journal of operational research*, 207(3):1293–1303.
- Suryanto, M., Plummer, V., Boyle, M., et al. (2017). EMS systems in lower-middle income countries: a literature review. *Prehospital and disaster medicine*, 32(1):64-70.
- The Ciw library developers. (2022). Ciwpython/ciw: v2.3.3.
- Toro-Díaz, H., Mayorga, M., Chanta, S., and McLay, L. (2013). Joint location and dispatching decisions for emergency medical services. *Computers & industrial engineering*, 64(4):917–928.
- Toro-Díaz, H., Mayorga, M., McLay, L., Rajagopalan, H., and Saydam, C. (2015). Reducing disparities in large-scale emergency medical service systems. *Journal of the Operational Research Society*, 66(7):1169–1181.
- Valenzuela, T., Roe, D., Nichol, G., Clark, L., Spaite, D., and Hardman, R. (2000). Outcomes of rapid defibrillation by security officers after cardiac arrest in casinos. *New England Journal of Medicine*, 343(17):1206–1209.
- Veser, A., Sieber, F., Groß, S., and Prückner, S. (2020). The demographic impact on the demand for emergency medical services in the urban and rural regions of bavaria, 2012–2032. *Journal of Public Health*, 23:181–188.
- Virtanen, P., Gommers, R., Oliphant, T., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S., Brett, M., Wilson, J., Millman, K., Mayorov, N., Nelson, A., Jones, E., Kern, R., Larson, E., Carey, C., Polat, İ., Feng, Y., Moore, E., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E., Harris, C., Archibald, A., Ribeiro, A., Pedregosa, F., van Mulbregt, P., and SciPy 1.0 Contributors (2020). SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272.
- Wang, W., Wu, S., Wang, S., Zhen, L., and Qu, X. (2021). Emergency facility location problems in logistics: Status and perspectives. *Transportation research part E: logistics and transportation review*, 154:102465.
- Yusvirazi, L., Ramlan, A., and Hou, P. (2018). State of emergency medicine in Indonesia. *Emergency Medicine Australasia*, 30(6):820–826.