Operations Research Perspectives 8 (2021) 100202

Contents lists available at ScienceDirect

Operations Research Perspectives

journal homepage: www.elsevier.com/locate/orp

Locating emergency vehicles: Modelling the substitutability of resources and the impact of delays in the arrival of assistance

José Nelas [a,c], Joana Dias [a,b,*]

[a] Faculdade de Economia, CeBER, Universidade de Coimbra, Av. Dias da Silva, 165, 3004-512 Coimbra, Portugal
[b] Inesc-Coimbra, Universidade de Coimbra, Rua Sílvio Lima, Pólo II, 3030-290 Coimbra, Portugal
[c] Centro Hospitalar e Universitário de Coimbra - Hospital Pediátrico, R. Dr. Afonso Romão, 3000-602 Coimbra, Portugal

ARTICLE INFO

Keywords:
Location
OR in health services
Emergency vehicles
Uncertainty

ABSTRACT

The quality and promptness of emergency assistance is highly dependant on the location of existing emergency vehicles. In this work, we propose a new model for optimizing emergency vehicles' location that takes into account the existence of different types of emergency vehicles and the level of care they can provide, the possibility of vehicles' substitution considering the hierarchy of levels of care and the explicit consideration of the progression of an emergency episode when the arrival of assistance suffers delays. The inherent uncertainty that exists in this problem is represented by a set of scenarios. A heuristic procedure for solving the problem was also developed. The model and algorithmic approach were tested using real data. It is possible to conclude that the application of stochastic location models that explicitly consider the evolution of the health condition of the victims when care is delayed can lead to better emergency coverage. The location of vehicles is indeed influenced by the explicit consideration of the impact of assistance time on the victims' conditions.

1. Introduction and literature review

Emergency Medical Systems (EMS) offer care that aims to mitigate the morbidity and mortality associated with sudden illnesses or injuries and guaranteeing that people are assisted in an effective and timely way. The assessment of EMS is a complex task, requiring a holistic analysis that takes into account the EMS connection with all other existing community health resources and its integration within the national health system [1]. Within this complex system, the location and organization of emergency care in a pre-hospital context is crucial to ensure a proper assistance.

1.1. The importance of the arrival time

One of the aspects that influences the success of pre-hospital assistance is the number of existing emergency vehicles and their location, since this significantly impacts the time it takes for assistance to arrive to emergency occurrences. This arrival time is a very important factor for the success of the health support provided since the health status of the victims of medical emergencies often tends to worsen over time if assistance is delayed. According to INEM (Portuguese National Institute of Medical Emergency) guidelines [2], the probability of survival of cardiac arrest victims with defibrillation rhythms decreases by around 10% to 12% for every minute without electrical defibrillation. If victims receive basic life support (BLS) during this waiting time this percentage decreases to 3% to 4%. Other works suggest that the arrival of support should not exceed 8 or 9 min for urban areas and 14 min for rural areas, reflecting the differences in emergency care accessibility for different geographic areas [3, 4]. Some countries benchmark assistance response times. For instance, the USA require the response time to be no more than 8:59 min for 90% of episodes in urban area. This time constraint is changed to 15 and 30 min for rural and wilderness areas, respectively [5]. In the United Kingdom, 75% of the most critical emergency calls must be attended within an 8 min time window and 95% of these episodes must receive assistance within 14 min in urban areas and 19 min in rural areas. Hong Kong has defined a 12 min limit for 92% to 95% of all cases in rural or urban areas, respectively.

The survival rate of the victims depends not only on the arrival time of the assistance, but also on the level of care that the victims need depending on the emergency situation that occurred. Some authors assume that the survival rate depends on whether the emergency was due to trauma or medical situations [6–12]. The assistance arrival time can influence directly the health condition of the victim, and can also have other indirect impacts, like the costs associated with longer recovery and

* Corresponding author at: Faculty of Economics, University of Coimbra: Universidade de Coimbra, Av.Dias da Silva, 165, 3004-512 Coimbra, Portugal.
E-mail address: joana@fe.uc.pt (J. Dias).

https://doi.org/10.1016/j.orp.2021.100202
Received 14 May 2021; Received in revised form 27 July 2021; Accepted 24 September 2021
Available online 28 September 2021

hospitalization times or the occurrence of secondary health situations resulting from this assistance delay [13, 14].

In Portugal, pre-hospital emergency care is organized according to the Integrated Medical Emergency System and is carried out with the collaboration of different entities, such as INEM, Voluntary Firefighters (VF), Professionals Firefighters (PF), Red Cross (RC), amongst others [2]. There are no established benchmarks for assistance times. However, INEM has published some performance indicators considering that assistance should arrive within 15 and 30 min of the emergency call for urban and rural areas, respectively [15]. According to these time limits, it would be possible to cover 75% and 91% of the urban and rural emergency episodes, respectively, considering only basic and immediate life support. When it is not possible to assure that the most adequate means are sent to the emergency occurrence, it is more important to assure that some assistance arrives as soon as possible than to wait too long until the most adequate vehicles are available. This sometimes means sending vehicles that are not the best ones considering the level of care they can provide and the level of care the emergency episode requires.

1.2. Brief literature review

The location of emergency vehicles is a prolific area of research, and there are many different approaches and different points of view that have been considered. The reader is referred to [16–19] for an overview of this area. We would like to highlight some recent works that have some connections with the work here presented. Table 1 summarizes the main characteristics of the cited works.

Peng et al. [23] describe two mathematical programming models, structured as two-stage stochastic programming models. In one of the models, the aim is to cover demand while minimizing costs. The other model explicitly considers the deterioration of the achieved coverage in more severe scenarios. The authors consider the possibility of reallocating vehicles according to their respective activity during the planning horizon defined. The model is applied to real data from Northern Ireland. Yoon et al. [24] prioritize care according to the different health conditions of patients. They also consider a two-stage stochastic model, and the existence of two different types of ambulances. A given degree of priority is associated with each patient. Probability-based travel times are considered. In this model, it is possible to consider the need of having more than one vehicle being sent to the same occurrence. However, the model does not represent the possibility of a patient getting worse while waiting for the arrival of assistance. Zhang and Zeng [25] present a two-stage stochastic model that considers the relocation of an ambulance to ensure the maximization of coverage determined by a given distance radius. There are reallocation costs that are distance dependant. With the objective of maximizing road accidents coverage, Mohri and Haghshenas [26] propose a network location model in which the vehicle's bases would be located at the most problematic areas, meaning that vehicles could be located at the edges instead of the nodes of the network. Time limits, average occurrence of accidents and other emergency events are considered. They do not require all the episodes to be covered and do not consider the possible unavailability of vehicles. Boujemaa et al. [27] consider an ambulance redeployment problem for two-tiered EMS systems that uses two types of ambulances for basic and advanced life support to respond to emergency calls with two priority levels. The proposed model considers time-dependant and stochastic demand in a multi-period ambulance redeployment setting. The authors develop two heuristic solution approaches and evaluate the performance of the model by simulation considering uncertainties in preparation and service time. The model assumes that only one ambulance is needed for each occurrence. A robust approach is proposed by Akıncılar and Akıncılar [28] with the objective of strengthening the choice of ambulance location against uncertainty in two dimensions: the average speed of the ambulance in reaching the accident site and the increase in the route distance if an alternative to the shortest path must be used due to traffic. The authors implicitly assume that the necessary ambulances are always available which can weaken the applicability of the model. Wajid et al. [29] present a model (Double Standard model) that maximizes double accident coverage in South Delhi. The authors conclude that it is possible to achieve complete coverage with fewer emergency vehicles than the existing ones. The data used represent traffic accidents only, and the ambulances are assumed to be always available. Intending to minimize the total response time for covering each episode, Bélanger et al. [30] describe a decision model considering the location of ambulances and the dispatching decisions. The authors use a recursive simulation-optimization approach to solve the problem, as this approach better approximates real situations.

Table 1
Recent works in this area of research.

Article	Two stage stochastic model?	Several vehicles in each occurrence?	Different types of vehicles?	Worsening of health status?	Objective function	Other characteristics
Peng, Delage & Li, (2020)	✓	×	×	×	Maximize coverage demand Reducing Costs	Allowed realocation of ambulances
(Yoon, Albert, & White, 2021)	✓	✓	✓	×	Maximize coverage, with penalization of non-covered calls	Prioritize calls according to the severity of episodes
(Zhang & Zeng, 2019)	✓	×	×	×	Maximize coverage under the normal situation and in the worst-case scenarios	Considers costs of reallocating ambulances
(Mohri & Haghshenas, 2021)	×	×	×	×	Maximize the coverage of facilities in each edge	Considers limits of time for each edge cover
(Boujemaa, Jebali, Hammami, & Ruiz, 2020)	✓	×	✓	×	Minimize total costs considering ambulance reallocation, dispatching, and unsatisfied coverage	Simulation technique used Allowed reallocation of ambulances
(Akıncılar & Akıncılar, 2019)	×	×	×	×	Minimizes the number of bases	Considers two dimensions of uncertainty: time and travel distance
(Wajid, Nezamuddin, & Unnikrishnan, 2020)	×	×	✓	×	Maximizes double coverage	Only considers traffic accidents
(Bélanger, Lanzarone, Nicoletta, Ruiz, & Soriano, 2020)	×	×	×	×	Minimize the total response time to satisfy the demand when each call is served by an ambulance on the dispatching list of the corresponding zone	Considers location and dispatching of ambulances
Our model	✓	✓	✓	✓	Maximization of the covered episodes as early as possible	Considers hierarchical substitutability of the vehicles and dispatching decisions

Although time is a dimension that, in an implicit or explicit way, must be present in EMS optimization location models, not all of the existing approaches include the impact of assistance time in the evolution of the emergency occurrence, acknowledging the influence that this time can have in the location decision-making process. Time has been used as a constraint in some models (for example [20–22]), considering deadlines before which the care should arrive. However, we are not aware of existing approaches in which the evolution of the needed level of care, consistent with the evolution of the health conditions of the victims, is explicitly considered. There are studies in the medical field that address the impact of the assistance delay in the victim's condition. However, there is a lack of location models that include simultaneously the effect of assistance time and possible delay with the differentiation of the level of assistance, as we are proposing in our approach. The change of the needed level of care influences the type of vehicles that should be deployed, which in turn can impact location decisions. This work aims at contributing to the current state-of-the-art in this area of research.

In this work, we present a model that considers, simultaneously, the existence of different vehicle types, capable of assuring different levels of care, and the possibility of one vehicle being substituted by another vehicle or set of vehicles that are equivalent in terms of the level of care they can provide. It is assumed that there are a set of potential and predetermined locations where the emergency vehicles can be located.

Moreover, it is also possible to explicitly consider the evolution of the emergency episode, by discretizing this evolution assuming that one episode can have different stages and establishing different needs for these different stages depending on the evolution of the victims' health conditions. The model also represents the assumption that it is better to have some assistance arriving, even if it is not the most suitable one, than not having any assistance at all.

Although the objective of the model is not to optimize the dispatching decisions of vehicles as emergency episodes occur, it is important to acknowledge the fact that the dynamic assignment of vehicles to episodes has impacts on the vehicles' availability and, inevitably, on the location decisions.

The contributions of this manuscript can thus be identified as:

- Developing a new model that includes, explicitly, the evolution of emergency episodes when assistance is delayed, an important feature that has not been considered so far.
- Allowing the evolution of the episodes to be discretized into a number of stages, as many as the decision-maker desires.
- Considering in an explicit way the impact that vehicles' dispatching decisions have on the location decisions.
- Representing in a more faithful way the decisions that are made in practice, assuming that it is better to send a less adequate vehicle than not sending any vehicle at all.
- Presenting a heuristic procedure to solve a two-stage stochastic problem of a high dimension in terms of number of variables and constraints.

This manuscript is organized as follows. In section 2 the model developed is presented. Section 3 describes the case study considered along with the available data and the methodology adopted. The heuristic procedure applied is also explained and the main computational results are discussed. Section 4 presents the main conclusions and possible future work. Fig. 1 depicts the structure of this manuscript and of the work developed.

2. Location model considering the evolution of emergency episodes

The occurrence of emergency episodes is inherently stochastic: it is not possible to anticipate where or when they will occur, or what will be the necessary means to be deployed. In the developed model, this uncertainty is represented by scenarios. The model is, thus, a two-stage stochastic model, where location decisions are made in the first stage and dispatching decisions are made in the second stage. Considering, in an explicit way, decision variables representing the dispatching decisions is very important because these decisions reflect the true availability of vehicles when a new emergency episode occurs.

The model assumes that all the episodes must receive some assistance, even if not the most adequate one. If there is no vehicle available or the ones available are not the most adequate ones within a given and defined time window after the beginning of an emergency occurrence (because they are not able to provide the type of care that is needed), then the episode will receive assistance from any available vehicle, regardless of the level of care it can provide. The model is also able to account for the evolution of the emergency occurrence where there are delays in the assistance: if the health status of the victims deteriorates, this can lead to changes in the necessary vehicles that need to be sent. The representation of this possibility brings the model closer to what happens in real life. The model does not intend to represent an optimal decision-making process for assigning vehicles to emergency episodes

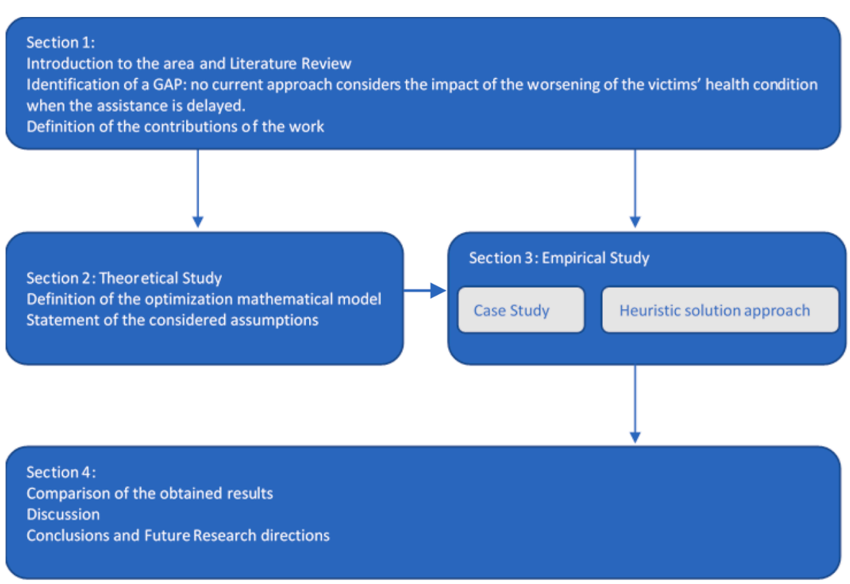

Fig. 1. Structure of this study.

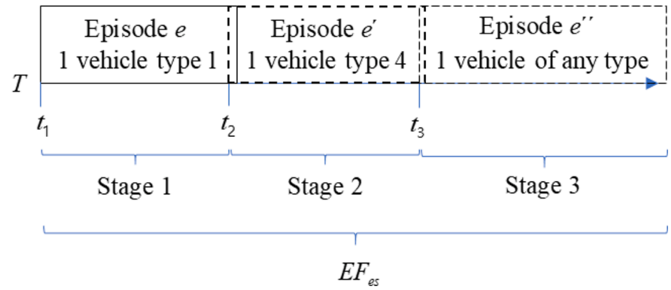

Fig. 2. Representation of an episode, through 3 fictitious episodes, considering delays in the arrival of assistance.

but should include this assignment problem in the most realistic way possible, since this assignment influences vehicle availability, and the corresponding location decisions.

Three levels of care will be considered (although the model is easily adapted if more levels of care are defined): Basic Life Support (BLS), Immediate Life Support (ILS) and Advanced Life Support (ALS).

Let us define the following notation:

T: time horizon

$i \in I$: set of possible bases for emergency vehicles

$j \in J$: set of locations where emergency episodes may occur

N_i: maximum number of vehicles that can be positioned in base i

$k \in K$: set of different types of vehicles, where each type k also determines the level of assistance of the respective vehicles ($l_k \in \{BLS, ILS, ALS\}, \forall k \in K$)

$v \in V$: set of existing vehicles, with each vehicle being characterized by a certain type (k_v)

p_s: probability of occurrence of each scenario $s, \forall s \in S$

$s \in S$: set of defined scenarios

$e \in E_s$: Set of real episodes that occur in a given scenario, $\forall s \in S$

Each emergency episode $e \in E_s$ is characterized by the location where it occurs, the number of vehicles of each level of care that should be sent to the location and the time periods in which the episode takes place. Considering the episode occurrence, the only time periods of interest are those related with vehicle assignments. When a vehicle is assigned to an episode it is not available for other emergency episodes that can occur during those time periods.

Each episode $e \in E_s$ is represented by n fictitious episodes that represent a discretization of the different stages of the episode's evolution. This is how the evolution of the episodes is represented in situations where it is not possible to immediately send the necessary emergency vehicles. Each episode $e \in E_s$ will thus be represented by a set of episodes EF_{es}, such that each fictitious episode e' belonging to EF_{es} represents the same real episode $e \in E_s$, but considering different time periods of occurrence. Episodes $e' \in EF_{es}$ may have different vehicle's needs reflecting, for example, the worsening of the victims' health state due to the late arrival of the means. It is also possible to consider situations in which vehicles are no longer necessary in later stages of the episode, because the victim was transported by other (private) means or, unfortunately, has died. Fig. 2 depicts this representation of emergency episodes.

episode if the assistance is delayed. The first stage (episode e) corresponds to the beginning of the emergency occurrence. This occurrence should receive one vehicle of type 1. If this vehicle is indeed timely assigned to the episode, then this episode ends, and it does not evolve to stages 2 and 3. However, if no vehicle arrives, then, at period t_2, there is a change in the evolution of this episode: instead of needing one vehicle of type 1, the episode now needs one vehicle of type 4. If such a vehicle is available and it is assigned to the episode, the episode ends when the vehicles assigned are released. If no vehicle is assigned at this stage, then a third stage is considered where, given the delay already confirmed in the assistance of this episode, the most important thing is to guarantee that some vehicle is sent, no matter its type. The evolution of this episode is, therefore, discretized into three different stages, which may consider different care needs. The last stage represents the assumption that it is better to send some assistance, even if not the most adequate one, than to send no assistance at all. These fictitious episodes can have intersections in their occurrence timings, and each episode can be represented by as many stages as one wishes.

Each episode e' will thus belong to one set EF_{es}, $e \in E_s$, and it will be characterized by:

$j_{e's} \in J$: the location where the episode takes place, $\forall s \in S$, that is the same for all the episodes belonging to EF_{es};

$n_{e'ks}$: number of vehicles of type k needed for that episode, $\forall s \in S, k \in K$

$TS_{e'k}$: the time period where assistance, considering vehicle type k, begins, $\forall k \in K$;

$TSt_{e'k}$: the time period where assistance, considering vehicle type k, ends, $\forall k \in K$.

For ease in the notation, EF_{es} is considered equal to $EF_{e's}$, $\forall e' \in EF_{es}$, meaning that set EF_{es} can be identified by any fictitious episode e' representing any stage of the real episode e (since each e' represents the stage of one, and only one, real episode).

It is assumed, with no loss of generality, that the set of all the episodes $e' \in \bigcup_{e \in E_s} EF_{es}$, are chronologically ordered, $\forall s \in S$.

It is possible to assign a binary value to each fictitious episode that will define whether that episode corresponds to a stage, within the episode's evolution, in which the best thing to do is to send any vehicle, whatever its type.

$$Q_{e's} = \begin{cases} 1, & \text{if episode } e' \in EF_{es} \text{ does not need to receive any type of vehicle in particular} \\ 0, & \text{otherwise} \end{cases}, \quad \forall e' \in EF_{es}, \forall e \in E_s, \forall s \in S$$

In the example depicted in Fig. 2, one episode is represented by three different fictitious episodes that correspond to the evolution of the

For each pair (i, j) and for each level of assistance l, the coverage matrix will give the information of whether j is in the coverage radius of base i or not. This coverage matrix represents the possibility of having different coverage time limits for different levels of needed assistance.

$$a_{ijl} = \begin{cases} 1, \text{ if base } i \text{ is within the coverage radius of } j \\ \text{defined for level of assistance } l \text{ (} j \text{ is covered by vehicles in } i, \text{ for assistance level } l\text{)}, \\ 0, \text{ otherwise} \end{cases} \forall i \in I, j \in J, l \in \{BLS, ILS, ALS\}$$

Notice that it is possible to have different maximum time limits defined for different levels of assistance, so that a_{ijl} can take different values for different values of l.

We also consider a substitution matrix, that makes explicit the hierarchical relationship that exists between vehicles that can be dispatched to a particular episode. A vehicle of type k' can replace another vehicle of type k if and only if it provides at least the level of care of type k (taking also into account transportation capability). This is a one-to-one substitutability (one vehicle substitutes another one).

$$c_{kk'} = \begin{cases} 1, \text{ if a vehicle of type } k' \text{ can substitute a vehicle of type } k \\ 0, \text{ otherwise}, \quad \forall k, k' \in K \end{cases}$$

$$z_{e's} = \begin{cases} 1, \text{ if } e' \text{ receives all the necessary vehicles} \\ \text{within the appropriate time interval, under scenario } s \\ 0, \text{ otherwise} \end{cases}, \forall e' \in \bigcup_{e \in E_s} EF_{es}, s \in S$$

It should be stressed that, in general, $c_{kk'}$ can be different from $c_{k'k}$ (k' can replace k but the opposite cannot occur). There are situations where a vehicle can be substituted by two or more additional vehicles. This is a many-to-one substitutability (more than one vehicle substitute one vehicle). These situations will be properly represented by constraints in the model, following [22].

Let us define:

$$r_k = \begin{cases} 1, \text{ if vehicle of type } k \text{ can be substituted by more than one vehicle} \\ 0, \text{ otherwise} \end{cases}, \forall k \in K$$

$g_k = \{(k', k'') :$ vehicle of type k can be substituted by a pair of vehicles k' and k'', $k', k'' \in K\}, k \in K$, assuming that $(k, k) \in g_k, \forall k \in K$.

To assure that the same vehicle is not assigned to episodes that have intersecting time periods, an incompatibility matrix is also built, defining whether two episodes are overlapping or not.

$$b_{e',e''ks} = \begin{cases} 1, \text{ if episode } e' \text{ and } e'' \text{ have intersecting time periods considering} \\ \text{the assistance by vehicles of type } k, \text{ so that a vehicle cannot} \\ \text{be simultaneously assigned to both, under scenario } s \\ 0, \text{ otherwise} \end{cases}, \forall e', e'' \in \bigcup_{e \in E_s} EF_{es}, \forall k \in K, s \in S$$

To consider the situation in which a given vehicle is not available for reasons other than being assigned to emergency cases, an availability matrix is also built.

$$d_{e'vs} = \begin{cases} 1, \text{ if vehicle } v \text{ can be assigned to episode } e', \text{ under scenario } s \\ 0, \text{ otherwise} \end{cases}, \forall e' \in \bigcup_{e \in E_s} EF_{es}, v \in V, s \in S$$

The decision variables are defined as follows:

$$y_i = \begin{cases} 1, \text{ if location (base) } i \text{ has vehicles located there} \\ 0, \text{ otherwise} \end{cases}, \forall i \in I$$

$$h_{vi} = \begin{cases} 1, \text{ if vehicle } v \text{ is located at } i \\ 0, \text{ otherwise} \end{cases}, \forall v \in V, i \in I$$

$$x_{ve'ks} = \begin{cases} 1, \text{if vehicle } v \text{ is assigned to } e', \\ \text{as being of type } k, \text{ in scenario } s \\ \text{(even if it is of a different type)} \\ 0, \text{ otherwise} \end{cases} \forall e' \in \bigcup_{e \in E_s} EF_{es}, v \in V, k \in K, s \in S$$

Each episode is represented by a set of fictitious episodes, that correspond to different stages in the episode's evolution. However, even considering these different stages and the possible delays in the arrival of assistance, it is not possible to assure that there are available and sufficient vehicles to assign to all the episodes, in one of the respective stages. Each stage has a beginning and ending time period, and the last stage in each episode is not extended until the end of the planning horizon considered. The possibility of having episodes that are not covered at all, due to lack of resources, is represented by the following binary decision variables:

$$adm_{es} = \begin{cases} 1, \text{ if real episode } e, \text{ belonging to scenario } s, \\ \text{does not receive any assistance} \quad , \forall e \in E_s, \forall s \in S \\ 0, \text{ otherwise} \end{cases}$$

The mathematical model is now presented. The objective function represents the maximization of the covered episodes, but considering weights associated with the different stages of each episode. These weights represent the decision-maker's preferences: it is better for any episode to receive assistance in the first stage than in a later one. So, weights $\varpi_{e's}, \forall e' \in EF_{es}, e \in E_s, s \in S$ are considered, such that $\varpi_{e's} >> \varpi_{e''s}$ for all $e', e'' \in EF_{es}$ with $e' < e'', \forall e \in E_s$.

In the second term of the objective function, M represents a very large positive number, meaning that solutions leading to episodes that do not receive any assistance at all are severely penalized.

$$Max\, Z = \sum_{s\in S} p_s \sum_{e\in E_s} \sum_{e'\in EF_{es}} \varpi_{e's} z_{e's} - M \sum_{s\in S} \sum_{e\in E_s} adm_{es} + \varepsilon\left(\sum_{v\in V}\sum_{i\in I} h_{vi} + \sum_{i\in I} y_i + \sum_{v\in V}\sum_{s\in S}\sum_{k\in K}\sum_{e\in E_s}\sum_{e'\in EF_{es}} p_s x_{ve'ks}\right) \quad (1)$$

Restrictions can be defined as follows:

- An episode is considered covered if and only it receives all the adequate vehicles, in terms of level of care. However, if $Q_{es} = 1$, then this constraint is redundant (in this case, the most important thing is to send any vehicle, regardless of the level of care it can provide, and constraints (8) apply). This constraint applies only if $r_k = 0$ (direct substitutability of one vehicle by the other is considered).

$$n_{eks} z_{es} \le \sum_{v\in V} c_{kk_v} x_{veks} + MQ_{es}, \forall s \in S, e \in \bigcup_{e'\in E_s} EF_{e's}, k \in K : r_k = 0 \quad (2)$$

- Constraints (3) to (7) do not apply if $Q_{es} = 1$ so, if this is the case, they are made redundant.

  Define auxiliary integer variables $O_{ekk's}, \forall s \in S, e \in \bigcup_{e'\in E_s} EF_{e's}, k \in K, k' : \exists(k', k'') \in g_k$ which represent the number of vehicles of type k' that are dispatched to episode e substituting vehicles of type k, under scenario s. Furthermore, let $q_{ekk'k''s}, \forall s \in S, e \in \bigcup_{e'\in E_s} EF_{e's}, k \in K : r_k = 1, (k', k'') \in g_k, k' < k''$, represent the number of vehicles belonging to the pair $(k', k'') \in g_k$ that are simultaneously dispatched to episode e under scenario s. Then, if vehicle k can be substituted by more than one vehicle (namely the pair $(k', k'') \in g_k$) the following constraints hold[1]:

$$n_{eks} z_{es} \le O_{ekks} + \sum_{(k',k'')\in g_k, k'<k''} q_{ekk'k''s} + MQ_{es}, \forall s \in S, e \in \bigcup_{e'\in E_s} EF_{e's}, k \in K : r_k = 1 \quad (3)$$

$$O_{ekks} \le \sum_{v\in V} c_{kk_v} x_{veks} + MQ_{es}, \forall s \in S, e \in \bigcup_{e'\in E_s} EF_{e's}, k \in K : r_k = 1 \quad (4)$$

$$O_{ekk's} \le \sum_{v\in V} c_{k'k_v} x_{vek's} - n_{ek's} + MQ_{es}, \forall s \in S, e \in \bigcup_{e'\in E_s} EF_{e's}, k, k' \in K : \exists(k', k'') \in g_k \quad (5)$$

$$\sum_{(k',k'')\in g_k : k'<k''} q_{ekk'k''s} \le O_{ekk's} + MQ_{es}, \forall s \in S, e \in \bigcup_{e'\in E_s} EF_{e's}, k, k' \in K : \exists(k', k'') \in g_k \quad (6)$$

$$\sum_{(k',k'')\in g_k : k'<k''} q_{ekk'k''s} \le O_{ekk''s} + MQ_{es}, \forall s \in S, e \in \bigcup_{e'\in E_s} EF_{e's}, k, k'' \in K : \exists(k', k'') \in g_k \quad (7)$$

- For episodes e such that $Q_{es} = 1$, the episode is considered covered if it receives at least one vehicle, whatever its type.

$$z_{es} \le \sum_{v\in V}\sum_{k\in K} x_{veks} + M(1 - Q_{es}), \forall s \in S, e \in \bigcup_{e'\in E_s} EF_{e's} \quad (8)$$

- All the real episodes should be covered in one, and exactly one, of the respective stages. When this is not possible for some episode, this situation is identified by having variable adm_{es} taking the value 1. If, when solving the problem, at least one adm_{es} is equal to 1 this means that the available vehicles are not sufficient to cover all the occurrences. To better understand the delays this lack of resources implies, one possibility is to increase the number of stages that represent the evolution of each episode (allowing for the existence of further

[1] Variable O_{ekks} represents the number of vehicles that are dispatched to episode e without being substituted by more than one vehicle. Episode e is covered if and only it receives the necessary vehicles, either vehicles that are of the required type or others that are equivalent considering either one-to-one or many-to-one substitutability. It is important to assure that no vehicle is counted more than once for the same episode. Additional details explaining these constraints can be found in [22].

delays and understanding in a more detailed way when will it be possible to send assistance).

$$\sum\nolimits_{e' \in EF_{es}} z_{e's} \geq 1 - Madm_{es}, \forall e \in E_s, s \in S \tag{9}$$

$$\sum\nolimits_{e' \in EF_{es}} z_{e's} \leq 1, \forall e \in E_s, s \in S \tag{10}$$

- An emergency vehicle can only contribute to a certain level of assistance in each episode of each scenario.

$$\sum\nolimits_{k \in K} x_{veks} \leq 1, \forall s \in S, e \in \bigcup\nolimits_{e' \in E_s} EF_{e's}, v \in V \tag{11}$$

- Emergency vehicle v can only assist episode e from i if it is located there, if e occurs within the coverage radius of i for that level of assistance, and if it is available in that scenario.

$$x_{veks} \leq d_{evs} \sum_{i \in I} a_{ij_e l_k} h_{vi}, \forall v \in V, s \in S, e \in \bigcup\nolimits_{e' \in E_s} EF_{e's}, k \in K \tag{12}$$

- An emergency vehicle can only be sent to two episodes if they occur in periods of time that do not intersect, considering the scenario to which both episodes belong.

$$\sum_{k' \in K} x_{vek's} + \sum_{k' \in K} x_{ve'k's} < 2 - b_{ee'ks'} \forall v \in V, k \in K, s \in S, e, e' \in \bigcup\nolimits_{e'' \in E_s} EF_{e''s} : e < e' \tag{13}$$

- There is a maximum number of vehicles that can be located at each base, and vehicles can only be assigned to a base that is prepared to receive them.

$$\sum_{v \in V} h_{vi} \leq N_i y_i, \forall i \in I \tag{14}$$

- Each emergency vehicle can only be assigned to one base.

$$\sum_{i \in I} h_{vi} \leq 1, \forall v \in V \tag{15}$$

- It is necessary to ensure that it is not possible to anticipate the future in each scenario, taking into account the decision to send emergency vehicles to occurrences. As, in the current model, each real episode is represented by a set of fictitious episodes these constraints should only be considered when no one of the fictitious episodes that corresponds to the same real episode is already being assisted.

$$\sum_{k' \in K} x_{ve'k's} \leq \left(2 - \sum_{i \in I} a_{ij_{e'} l_k} h_{vi} - \sum_{i \in I} a_{ij_e l_k} h_{vi}\right) + 1 - \left(b_{ee'ks} - x_{veks} - \sum_{e'' \in EF_{es}} z_{e''s}\right) + \frac{\sum_{v' \in V} x_{v'eks}}{n_{eks}} + M(1 - c_{kk_v}) + \sum_{k' \in K} \sum_{e'' : e'' < e} b_{e''ek's} x_{ve''k's}, \forall v \in V, s \in S, e, e' \in \bigcup_{e'' \in E_s} EF_{e''s} : e < e', k \in K : n_{eks} \geq 1 \tag{16}$$

The behaviour of the variables can be defined as follows:

$y_i \in \{0, 1\}, \forall i \in I$
$h_{vi} \in \{0, 1\}, \forall v \in V, i \in I$
$adm_{es} \in \{0, 1\}, \forall e \in E_s, s \in S$
$z_{es} \in \{0, 1\}, \forall e \in \bigcup_{e' \in E_s} EF_{e's}, s \in S$
$x_{veks} \in \{0, 1\}, \forall v \in V, e \in \bigcup_{e' \in E_s} EF_{e's}, k \in K, s \in S$
$q_{ekk'k''s} \geq 0$ and integer, $\forall e \in \bigcup_{e' \in E_s} EF_{e's}, s \in S, k, k', k'' \in K : (k', k'') \in g_k$
$q_{ekk's} \geq 0$ and integer, $\forall e \in \bigcup_{e' \in E_s} EF_{e's}, s \in S, k, k' \in K : (k', k'') \in g_k$

3. Case study

3.1. Available data

The case study considers the emergency episodes that occurred in 2017, in the district of Coimbra, Portugal. All the data was totally anonymized and provided by INEM. In this civil year, a total of 50,732 emergency episodes occurred, requiring 60,343 vehicles' dispatches. The difference between the number of vehicles needed and the number of occurrences means that there were occurrences needing more than one vehicle (because there are several victims or because the needed care can only be provided by using more than one vehicle, for instance).

There are five types of vehicles that provide three different levels of assistance: BLS, ILS and ALS.

BLS vehicles can be Medical Emergency Motorcycles (MEM), Assistance Ambulances (AA) or Medical Emergency Ambulances (MEA). MEM are vehicles manned by an emergency technician able to deliver BLS and External Automatic Defibrillation (EAD). AA are manned by volunteers or professionals with specific training in prehospital emergency techniques. They can work with other means of emergency and are able to transport the victims to health units. MEA are manned by two emergency technicians and their mission is to stabilize the victim autonomously or in complementarity with other means, and to transport the victim to the hospital. They have equipment for resuscitation and clinical stabilization, namely EAD.

Immediate Life Support Ambulances (ILSA) are ILS vehicles that guarantee more differentiated health care than the previous means, such as resuscitation maneuvers. Their crew consists of a nurse and a prehospital emergency technician. They can work in partnership with other means (with different levels of care differentiation) and have the capacity to conduct defibrillation, cardiac monitoring and transmit electrocardiographic data. They can also transport victims.

Medical Emergency and Resuscitation Vehicles (MERV) are ALS vehicles that are used for prehospital intervention. Their crew counts with a nurse and a medical doctor with competence and equipment for ALS. They aim to stabilize and monitor the transportation of victims to the hospital. These vehicles do not have the ability to transport victims, which forces the dispatchment of another emergency vehicle whenever there is such a need.

For each real episode, it is possible to know the place where the episode occurred and the vehicles that were dispatched, as well as the dispatching times. One important information that is missing is the exact

Table 2
Substitutability matrix (values for $c_{kk'}$).

K	1	2	3	4	5
1	1	0	1	1	0
2	1	1	1	1	0
3	1	0	1	1	0
4	0	0	0	1	0
5	0	0	0	0	1

Table 3
Total number of available vehicles, per type and base.

Base	k					Base	k				
	1	2	3	4	5		1	2	3	4	5
Coimbra VF	1	0	0	0	0	Mira RC	0	0	0	0	0
Coimbra PF	1	0	0	0	0	Mira VF	1	0	0	0	0
Coimbra RC	1	0	0	0	0	Cantanhede VF	1	0	0	0	0
Coimbra Central Hospital (CH)	0	0	0	0	1	Penela VF	1	0	0	0	0
Coimbra General Hospital (GH)	0	0	1	0	1	Miranda do Corvo VF	1	0	0	0	0
Borda do Campo RC	1	0	0	0	0	Lousã VF	1	0	0	0	0
Pereira RC	1	0	0	0	0	V. N. de Poiares VF	1	0	0	0	0
INEM Regional Base (RB)	0	1	2	0	0	Penacova VF	1	0	0	0	0
Condeixa VF	1	0	0	0	0	Tábua VF	1	0	0	0	0
Soure VF	1	0	0	0	0	Góis VF	1	0	0	0	0
Cantanhede Basic Urgency (BU)	0	0	0	1	0	Arganil BU	0	0	0	1	0
Montemor-o-Velho VF	1	0	0	0	0	Arganil VF	1	0	0	0	0
Figueira da Foz PF	0	0	0	0	0	Pampilhosa Serra VF	1	0	0	0	0
Figueira da Foz VF	1	0	0	0	0	Oliveira de Hospital VF	1	0	0	0	0
Figueira da Foz DH	0	0	1	0	1	Serpins VF	1	0	0	0	0
Figueira da Foz RC	1	0	0	0	0	Coja VF	1	0	0	0	0
Carvalhais RC	1	0	0	0	0	V. N. de Oliveirinha VF	1	0	0	0	0
Maiorca RC	1	0	0	0	0	Laborins RC	1	0	0	0	0
Quiaios RC	1	0	0	0	0	Lagares da Beira VF	1	0	0	0	0
Carapinheira RC	1	0	0	0	0	Brasfemes VF	1	0	0	0	0
Verride RC	1	0	0	0	0						

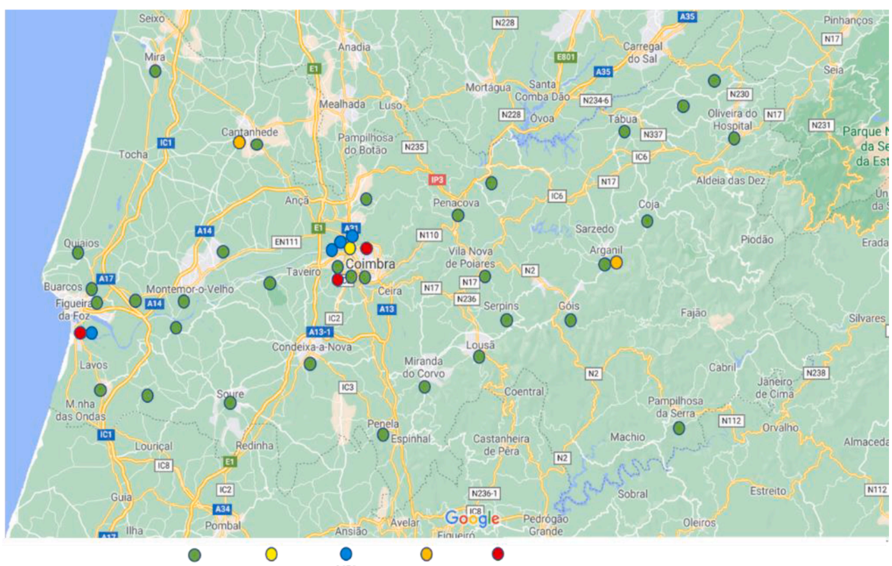

Fig. 3. Geographical distribution of the existing vehicles in 2017.

time each occurrence started (the exact time when the emergency call was received). The information regarding the total assignment time of a given vehicle to a given episode is also missing (only departure time is known, and not the time when each vehicle is operational again). As we have already worked with other datasets for the same region, we have considered similar assignment times. It is assumed that the total duration (in minutes) of each vehicle assignment follows a Gamma distribution, with shape α=4.2123 and rate β=0.0692 [22].

There are a total of 43 available vehicles, that are distributed by 41 bases, such that:

- 33 vehicles are of the type AA represented by $k = 1$
- 4 vehicles are of the type MEA, $k = 2$
- 1 vehicle is of type MEM, $k = 3$
- 2 vehicles are of the type ILSA, $k = 4$
- 3 vehicles are of the type MERV, $k = 5$

These existing 41 bases are the ones considered as being the possible locations for the vehicles (this set constitutes set I).

Table 2 shows the direct substitutability that is possible to be done, regarding vehicle's types (one-to-one substitutability). Moreover, it is considered that a vehicle of type 4 can also be substituted by sending simultaneously a vehicle of type 5 and a vehicle of type 1 or 3. This means that $r_4 = 1$ and $g_4 = \{(4,4), (1,5), (3,5)\}$

Table 3 shows all the emergency vehicles that are available by type of vehicle. The current location of the vehicles, defined in Table 3, is also depicted in Fig. 3. This solution is called the "current solution".

It is possible to observe that the most populated municipalities, that also correspond to the municipalities with a greater incidence of emergency occurrences (Fig. 4), are the ones that count with a greater number of vehicles. There are ILSA vehicles in Arganil and Cantanhede municipalities mostly due to the fact that only basic medical urgency services are located in these areas covering also other municipalities (such as Mira, Tocha, Tábua, Góis, Vila Nova de Poiares, Oliveira do Hospital, Pampilhosa da Serra, Lousã, Miranda do Corvo).

Considering the level of care needed, 87.85% of the emergency

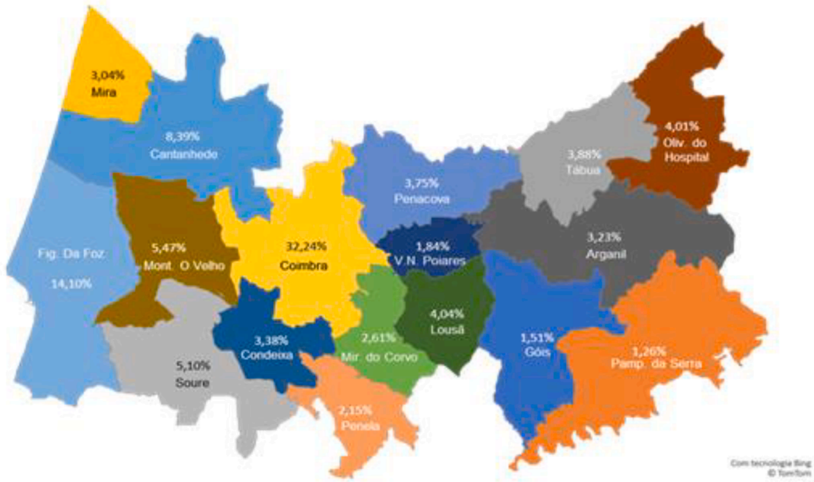

Fig. 4. Distribution of the emergency occurrences by municipality.

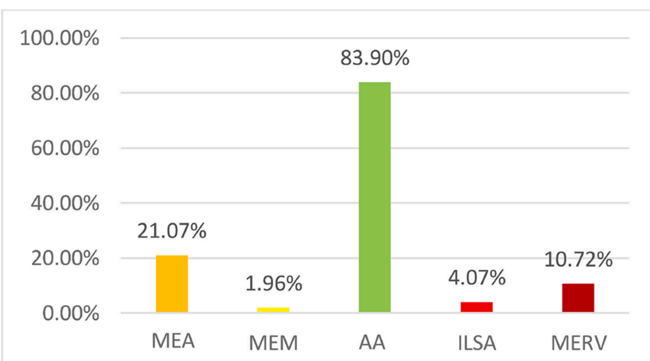

Fig. 5. Distribution of the dispatchment of vehicles considering the vehicle type.

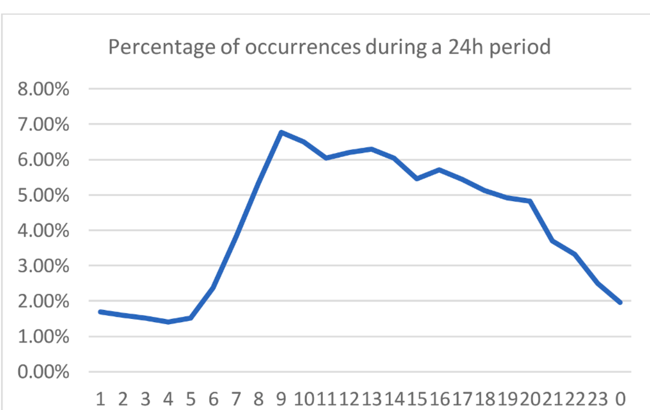

Fig. 6. Evolution of the percentage of occurrences in a 24-hour period.

episodes required an assistance of level BSL, 8.86% ALS and the rest ILS.

Fig. 5 depicts the frequency with which each type of vehicle was sent to the emergency occurrences.

The most used vehicle type is AA, that is sent to 83,90% of the episodes. MEA vehicles are sent to 21,07% of the episodes. The single MEM vehicle available is sent to only 1,96% of the cases. This is probably explained by the fact that only one vehicle exists, it has no transportation capability, and it can be easily replaced by other types of vehicles (MEA, ILSA or even AA). Vehicles of types ILSA and MERV, the ones that have the most differentiated level of care, are sent in 4,07% and 10,72% of the times, respectively.

The average number of occurrences changes during the day, as can be observed in Fig. 6. It is possible to notice the decrease in the number of episodes in the 22 h - 6 h period, when compared with the rest of the day.

In this work, we have not studied in detail the effects of seasonality in the occurrence of emergency episodes. There are some municipalities where this seasonality is more pronounced, namely during August. In municipalities near the coast, this is explained due to the affluence of tourists. In other municipalities this increase can reflect flows associated with emigrants. Fig. 7 depicts the percentage of the episodes that occur in four municipalities during each month of the year. These municipalities were chosen because they are the ones in which this seasonality is more pronounced. Figueira da Foz and Mira clearly present more episodes in August (coastal municipalities), as well as Góis (due to the arrival of emigrants). Coimbra presents the opposite behaviour since many of its inhabitants leave this area during this month. Some increase can also be observed near Christmas and New Year's Eve.

In the future, it can be useful to study alternative locations for emergency vehicles during the months of August and December.

3.2. Methodology

The proposed model was applied to the case study described. The objective is to calculate a new solution, considering the available data, and to assess if the calculated solution is similar or different from the current solution. If the calculated solution is different, then both solutions should be compared.

The methodology considered was structured into three different steps. In the first step, 100 days from the 365 days of data available (the whole civil year of 2017) where randomly selected, so that the corresponding data was used as input data for the model (one day will represent one scenario in the model). Then, a solution was calculated (second step). In the third step, both the current and the calculated solutions were tested considering not only this first dataset, but also another dataset corresponding to 50 days, also randomly selected from the 365 days available. These sets (100 samples and 50 samples) are disjoint, but they are both randomly generated from the real dataset. This third phase can be interpreted as an out-of-sample assessment: the calculated solution is assessed on a completely different dataset than the one that was used by the model and that generated that solution. This allows a fairer comparison between the current and the calculated solutions. Otherwise, the comparison could be biased, benefiting the calculated solution since it would be assessed based on the same dataset

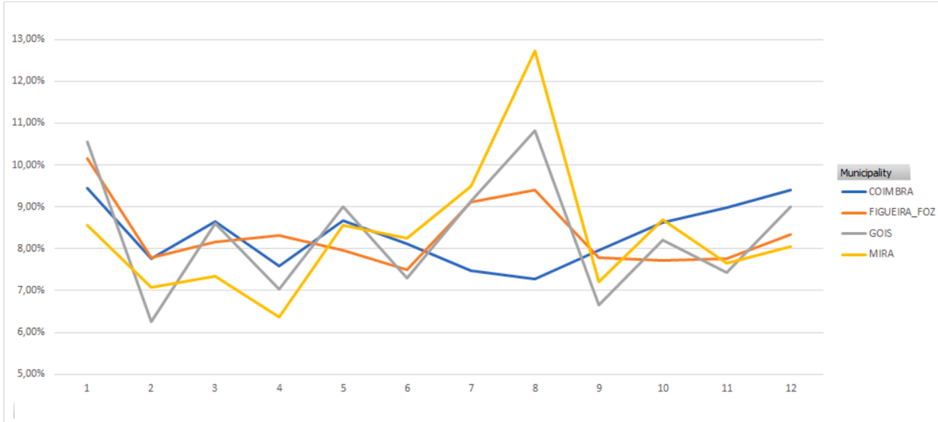

Fig. 7. Percentage of episodes that occur, in each municipality, in each month (2017 year).

that originated it.

3.3. In-sample set

To generate the data that will be the input for the model, 100 randomly generated days are selected from the 2017 real database. Each day will constitute one scenario. In these randomly selected 100 days, around 13,700 emergency episodes have occurred.

In Portugal, each municipality is organized into smaller administrative regions called *freguesias*. We considered a total of 209 *freguesias*. Instead of considering the exact GPS coordinates where each emergency episode occurred, all the episodes that occurred in one *freguesia* are considered as taking place in the *freguesia*'s centre. As *freguesias* have a small and limited geographical area, this simplification, that speeds the calculation of the driving times (as explained later) does not introduce significant errors that could impact the calculated solution.

From these randomly selected days, it is possible to know exactly how many episodes occurred, and where they occurred (the *freguesia* where they occurred). Moreover, it is possible to know which were the vehicles that were dispatched and when they were sent. We do not have access to the exact total time these vehicles were assigned to these episodes. So, for each episode that occurred in each selected day, the total duration of the episode is randomly generated, as explained in [22].

Two assumptions were made: 1) as we do not know the timing of the emergency call (the exact timing where the emergency episode really began), we will assume that the first vehicle sent to the episode corresponds to the beginning of the episode; 2) we will assume that the vehicles sent are exactly the ones that should have been sent, so these vehicles will implicitly represent the level of care and the number of vehicles that the episode needs.

The evolution of these episodes is represented, initially, by three fictitious episodes: we are discretizing the evolution of each episode by considering three different stages. The starting time of one stage is displaced 10 min from the start of the previous one. In the third stage, we allow any vehicle, independently of the level of care it can provide, to assist the emergency episode. The total timing of this last stage is randomly generated by multiplying the initially randomly generated duration by a number in the interval [1, 1.5], considering a uniform distribution. For each episode, the possibility of the status of the victims worsening from one stage to the next, leading to a change in the vehicles' needs, is randomly generated. This random generation takes into consideration the analysis of the existing data: by analysing each occurrence and the sequence in which different vehicles where sent (if this was the case), as well as the main cause of the occurrence, it is possible to generate the probability of a given episode evolving unfavourably, changing the needed vehicles accordingly.

Table 4
Illustrative example of the heuristic procedure: optimal location of the vehicles for each scenario in iteration 1.

	Scenarios				
Vehicles (type)	1	2	3	4	5
1 (1)	A	B	A	B	B
2 (1)	A	A	A	B	A
3 (1)	B	B	A	B	A
4 (2)	B	B	A	A	A
5 (2)	B	B	B	A	B

Table 5
Illustrative example of the heuristic procedure: optimal location of the vehicles for each scenario in iteration 2.

	Scenarios				
Vehicles (type)	1	2	3	4	5
1 (1)	*A*	*A*	*A*	*A*	*A*
2 (1)	*B*	*B*	*B*	*B*	*B*
3 (1)	A	B	B	A	A
4 (2)	*B*	*B*	*B*	*B*	*B*
5 (2)	B	B	A	A	B

For the coverage matrix, all the distances between the *freguesia* centres and all the existing bases were calculated using Google Maps, and considering driving times. A coverage time limit of 15 min was considered for urban areas and 30 min for rural areas.

If the instance is solved and there are episodes not being covered at all ($adm_{es} = 1$, for some episode $e \in E_s$), then the discretization of the evolution of all the episodes in the corresponding scenario will consider an additional stage (the fourth), having the same characteristics as the third, but starting 10 min later.

Regarding location variables, it was decided to fix all the existing AA vehicles in their current locations. This decision is justified by the resistance that exists in changing the location of the vehicles of this type [22]. Most of the AA vehicles are placed in firefighter stations, that are not willing to let these vehicles go anywhere else. These vehicles are a guarantee of INEM support to the institutions, in the form of staff training, access to differentiated equipment, as well as financial support for the vehicles' costs during the first four years of operation, and an additional support for every vehicle activation. Having INEM AA vehicles in the firefighter station is also seen as giving visibility and credibility to the service provided, being an important motivational factor to promote the attraction of volunteers and professionals. Furthermore, it promotes a sense of safety to local populations. Some of these vehicles belong to volunteer firefighter stations, or to the Red Cross, and many have been offered to these institutions by local companies which would not accept to have these vehicles transferred to other locations. These institutions are operated resorting almost fully to volunteer work, so any

Table 6
Comparison between the calculated and the current solution.

Vehicle	Current Solution	Calculated Solution
K	Base	Base
3	Figueira da Foz DH	Mira VF
3	INEM RB	Tábua VF
3	INEM RB	Miranda do Corvo VF
3	Coimbra GH	BV Brasfemes
4	Cantanhede BU	INEM RB
4	Arganil BU	INEM RB
5	Coimbra CH	Coimbra PF
5	Coimbra GH	INEM RB

change in the location of the means could jeopardize the continuity of their operations.

The weights that need to be defined in the objective function were set to 1000 for episodes receiving assistance in the first stage, 100 for episodes receiving assistance in the second stage, and 0 for all other situations.

3.4. Calculating a solution

As each real episode is represented by, at least, three fictitious episodes (there can be situations where this number needs to be higher, namely if adm_{es} is equal to one for some $e \in E_s$), the total number of episodes to consider in the model is huge, surpassing 41,100. The model dimension is not compatible with the use of a general solver. Cplex was not able to solve the instance created, due to memory issues. Computational results showed that it is only capable of solving instances considering, at most, three scenarios. Considering this computational limitation, a heuristic procedure was devised to calculate a solution.

Instead of creating one single instance incorporating 100 scenarios, 100 instances were created, each one considering only one scenario. These instances/scenarios are interpreted as different "experts" that, considering the limited information each of them has, propose a given location solution. The 100 instances are solved, and the corresponding optimal solutions are analysed. Each instance is solved in less than one minute of computational time, considering an Intel Xeon Silver 4116, 2.1 gigahertz, 12-core processor, 128 gigabyte RAM computer, and Cplex 12.7.

If there are location decisions with which more than 50% of the experts agree with, then the corresponding location variables are fixed. If there are still vehicles that are not yet located, the 100 instances are solved again, now considering a set of already fixed location variables, and the process is repeated until all the decisions considering the location of the vehicles have been determined. It is important to note that it is not necessary that the location variables are, themselves, exactly equal for the majority of the experts. What is relevant is the type of vehicle that is being located at each base. For instance, if vehicle 1 and 2 are of the same type k, expert 1 places vehicle 1 in base 2 and expert 2 places vehicle 2 in base 2, they are both agreeing that one vehicle of type k should be placed in base 2.

If, in some iteration, there is no location decision with which more than 50% of the experts agree, then the location decision receiving the greater number of votes is fixed (and only this one). The procedure stops when all the location variables have been fixed. The heuristic procedure is now presented.

1 Initialize $F = \emptyset$.
2 Solve all the existing instances, each one representing one scenario only, and fixing all variables h_{vi} such that $(v, i) \in F$ equal to 1 in all these instances.
3 Select all $k \in K$ such that it was decided to place at least one vehicle v of type $k_v = k$ in base i in more than 50% of the instances. For each one of these types k, choose one vehicle v such that $k_v = k$ and there is no pair $(v, i') \in F$. Fix $h_{vi} = 1, F = F \cup \{(v, i)\}$.
4 If there are still vehicles v such that there is no pair $(v, i) \in F$, go to step 2. Else stop. All the location variables have been fixed.
5 If no location variable was fixed in step 3, then choose one variable h_{vi} such that the choice of placing one vehicle of type k_v in base i is the most voted option amongst all the different location decisions calculated and there is no pair $(v, i') \in F$. Fix $h_{vi} = 1, F = F \cup \{(v, i)\}$. If there are still vehicles v such that there is no pair $(v, i') \in F$, go to step 2. Else stop. All the location variables have been fixed.

Let us consider a simple and small example to illustrate this heuristic procedure. Assume that there are 3 possible bases for locating vehicles (A, B, C), there are 5 vehicles that can be of two different types. Vehicles 1 to 3 are of type 1 and vehicles 4 and 5 are of type 2. Five scenarios are generated. Following the heuristic procedure, $F = \emptyset$ and all the five instances, one for each scenario, are solved. The optimal location solutions for each one of the scenarios are represented in Table 4.

Looking at Table 4, it is possible to observe that, in 4 out of 5

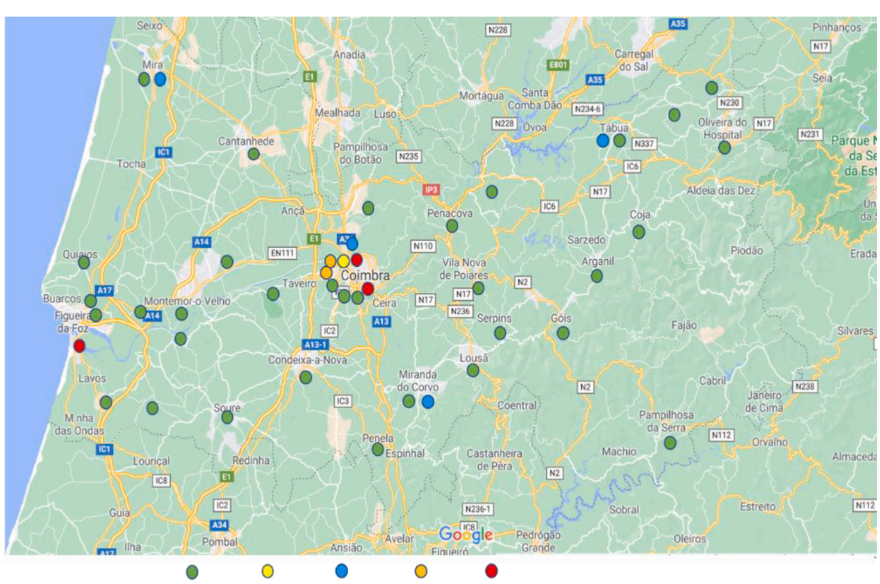

Fig. 8. Distribution of emergency vehicles proposed by the model.

Table 7
Coverage obtained with the current and calculated solutions for in-sample data.

Episode stages	Current solution				Calculated Solution			
	Average	Greatest	Smallest	Number of episodes	Average	Greatest	Smallest	Number of episodes
1	97.01%	100.00%	92.24%	13,086	95.26%	99.24%	87.72%	13,322
2	0.87%	3.20%	0.00%	154	1.15%	4.62%	0.00%	116
3	1.84%	5.43%	0.00%	451	3.60%	7.89%	0.74%	288
Without coverage (stage 4)	0.28%	0.00%	0.88%	35	0%	0%	0%	0

Table 8
Coverage obtained with the current and calculated solutions for out-of-sample data.

Episode stages	Current solution				Calculated Solution			
	Average	Greatest	Smallest	Number of episodes	Average	Greatest	Smallest	Number of episodes
1	95,45%	99,24%	89,66%	6556	96,67%	100%	92,24%	6640
2	0,92%	2,67%	0%	60	0,93%	3,20%	0%	62
3	3,36%	6,98%	0,74%	228	2,14%	5,43%	0%	144
Without coverage (stage 4)	0,28%	0,88%	0%	20	0,25%	2,71%	0%	18

scenarios, a vehicle of type 1 should be located at base A. As this represents more than half of the total number of scenarios then, according to step 3, one vehicle of type 1 is chosen, and its location is fixed at base A. As there are three vehicles of type 1, any of the three can be chosen. Let us consider vehicle 1: $h_{1A} = 1, F = F \cup \{(1, A)\}$.

In 3 out of 5 scenarios, another vehicle of type 1 is also located at base A. So: $h_{2B} = 1, F = F \cup \{(2, B)\}$.

Considering now vehicles of type 2, one vehicle of this type is located at base B in the majority of the scenarios, meaning that $h_{4B} = 1, F = F \cup \{(4, B)\}$.

As there are still vehicles v such that there is no pair $(v, i) \in F$, namely vehicles 3 and 5, we should solve the five instances again, but now considering some variables fixed, namely $h_{1A} = 1, h_{2B} = 1,\ h_{4B} = 1$. Table 5 presents the optimal solutions obtained in this second iteration of the procedure.

Performing a similar analysis, the procedure would now fix the location of vehicle 3 at base A and vehicle 5 at location B. This would end the procedure, since all the location decisions are determined.

Step 5 of the procedure accounts for the possibility of no single location decision gathering the votes of at least 50% of the experts (scenarios). If this is the case, the option most voted is the chosen one, and the process continues.

This heuristic procedure is not capable of guaranteeing the calculation of an optimal solution. However, it is possible to calculate an upper bound for the optimality gap. Actually, in the first iteration, the optimal solution for each one of the instances (scenarios) is calculated. This means that each one of these solutions considers the optimal location of vehicles for each specific scenario. Let Z_{heur} represent the objective function value of the solution obtained by the heuristic, Z^* the optimal objective function value of (1) and let $Z_s, \forall s$, represent the optimal objective function value for the instance that represents scenario s only. Then

$Z_{heur} \leq Z^* \leq \sum_{s \in S} p_s Z_s$, and an upper bound for the optimality gap when Z^* is not known can be calculated as $gap \leq \frac{\sum_{s \in S} p_s Z_s - Z_{heur}}{\sum_{s \in S} p_s Z_s}$.

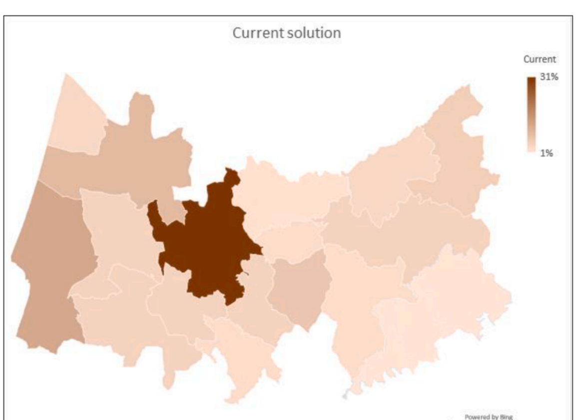

Fig. 9. Percentage of episodes assisted in stages 3 or 4 in each municipality with the current solution.

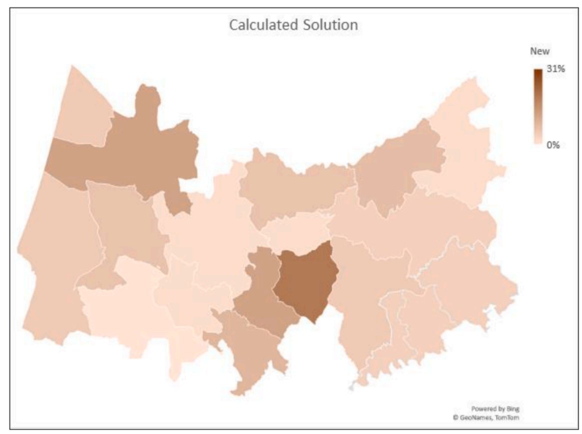

Fig. 10. Percentage of episodes assisted in stages 3 or 4 in each municipality with the calculated solution.

Other approaches were tested, namely using metaheuristics, but the results obtained were worse than this heuristic approach, both considering computational time and solution quality. One of the difficulties of applying a metaheuristic to this problem has to do with the huge number of constraints that have to be satisfied, and that guarantee the correct deployment of the vehicles. Finding a proper solution representation that does not promote the creation of unfeasible solutions is very difficult. Matheuristics have also been tried, where only the location variables were represented, and the second stage problems were solved by a general solver. The results were also much worse than the ones obtained with this heuristic approach.

For the data considered, it was always possible to fix location decisions based on the majority of more than 50% of the votes (it was not necessary to resort to step 5). The resulting optimality gap is less than 1.92%.

Table 6 compares the location decisions obtained (calculated solution) with the locations of the vehicles in the current solution. In Fig. 8 we can observe the changes in the geographic distribution proposed by our model, in comparison with the situation depicted in Fig. 3. We can see that, in the calculated solution, ILSA vehicles are relocated from two rural areas (Arganil and Cantanhede) to Coimbra urban area, and two out of three MEA vehicles are changed from urban to rural areas (from Coimbra to Tábua and Miranda do Corvo).

In the next section we analyse and discuss the obtained results.

4. Discussion of the results, conclusions and future work

4.1. Comparing the current and calculated solutions

The analysis of Table 6 shows important changes in the location of most of the considered vehicles. The location of four vehicles of type 3 is changed, and two ILSA vehicles are also relocated. There is a clear reinforcement of vehicles in the most inner parts of the district, in areas that are classified as rural. This seems to be a more balanced solution for emergency episodes that require ALS vehicles, and it can probably address some accessibility issues in less populated areas. Furthermore, the location of vehicles with more differentiated levels of care can have an indirect positive impact in the quality of care provided, since it can leverage the training of other human resources in the emergency assistance chain, especially those assigned to AA vehicles.

In an opposite direction, it is possible to observe the change in the location of ILSA vehicles, that are now placed in the heart of Coimbra city. As we can observe by the coverage results obtained, these decisions are actually balancing the reinforcement of the assistance to more rural areas, making essential differentiated assistance available where most of the population resides.

The location of one MERV in INEM RB can be justified by the fact that, in this location, the vehicle is almost equidistant from the city centre and two important communication roads with heavy traffic (highway n° 1 and IP3). This alternative is also compensating for the fact that one ILSA is withdrawn from Cantanhede, since this vehicle will be nearby. The choice of locating another MERV in Coimbra PF makes this vehicle closer to the North-eastern area, improving the coverage of the municipalities located there.

The current and calculated solutions are compared considering both the in-sample and out-of-sample datasets. The in-sample set considers 13,726 real episodes (corresponding to 41,178 fictitious episodes). The out-of-sample set is composed by 6864 episodes, corresponding initially to 20,592 fictitious episodes.

The coverage results obtained with the in-sample set are presented in Table 7. This table shows the total number of episodes that were covered in each of the stages considered. The average, lower and greater values presented considered the analysis of the daily coverage, for the 100 days selected.

For the calculated solution, it was possible to cover 100% of the emergency occurrences considering each real episode represented by three fictitious episodes. However, when the current solution was applied to this dataset, it is possible to see that 35 real episodes were not covered during the three stages. To tackle this situation, for all the days where this situation occurred, each episode was represented by four, instead of three, stages, so that a solution could indeed be found. It was sufficient to consider this additional time stage. Episodes that need this fourth stage are considered as not being properly covered.

From Table 7 it is possible to observe that both solutions are capable of achieving very good coverage results, since most episodes receive the adequate assistance in the adequate time window. The current solution is capable of covering in an optimal way, on average, 97,01% of the daily occurrences. This value is lower in the calculated solution (95,26%). However, the number of total episodes covered in the first phase is higher with the calculated solution. There are more episodes covered in stage 2 or 3 with the current solution than with the calculated solution.

As the results obtained with the in-sample dataset could be biased, and could benefit the calculated solution, these solutions were also compared with the out-of-sample dataset. The results obtained are summarized in Table 8. Once again, the calculated solution is able to cover, in the first stage, a higher number of episodes. The worst results considering daily data are also better: 92.24% of the episodes are covered in the worst case with the calculated solution, whilst 89.66% are covered with the current solution. Whereas the worst relative result considering the episodes that are not covered in stages 1 to 3 is poorer for the calculated solution than for the current solution (2.71% versus 0.88%), this corresponds to a lower number of episodes (18 versus 20). Achieving better results for the first stage of assistance is very important, since it can prevent the deterioration of the health state of the victims.

The change in the location of ILSA, that is currently located at Arganil, has an impact on the increase of the episodes that are not covered in the first stage in this municipality. This change makes the ALS vehicles further away, and they are most likely to arrive only at stages 2 or 3. Regarding the change in the location of vehicles that belonged to the area of Cantanhede municipality, it is possible to see that the relocation of these vehicles did not jeopardize the assistance in this area. The relocation of three BLS vehicles and 2 ILSA vehicles contributed to a best overall coverage of the more serious episodes, which need ALS.

It is also interesting to see which solution can be interpreted as being more equitable considering the different municipalities under study. As we want to prevent emergency episodes to be assisted in stages 3 or higher, we have considered the regional distribution of the total episodes assisted in these later stages. Fig. 9 and Fig. 10 depict this comparison, considering the distribution, in percentage, of these total number of episodes worst covered by municipalities. As can be seen, the calculated solution presents a smoother distribution of worst covered episodes. In the current solution, one of the municipalities with a larger number of worst covered episodes was Coimbra. It is possible to see a clear improvement with the calculated solution, due to the more reasonable location of means in urban areas.

4.2. Main conclusions and contributions

In this work, a new model for emergency vehicle location is presented that explicitly takes into account the worsening of the victims' state due to delays in assistance time, by representing the evolution of each episode by a discrete set of fictitious episodes. This is a very interesting and important feature since it mirrors more accurately real emergency assistance scenarios. This work focused on the impact the episode's evolution has in the types of vehicles that should be sent to the occurrence. However, the same type of reasoning can be used to represent temporal changes in other episode's features. This representation of each episode also allows for a better characterization of the quality of care provided and of the timing until assistance arrives. Furthermore, this model allows for the explicit consideration of vehicles associated with different levels of care, also including the possibility of vehicle substitutability.

The model was applied to a case study, resorting to real data. It is possible to conclude that the calculated solution assures a better coverage of the episodes in an earlier period of distress than the current solution. This can be observed in either in-sample and out-of-sample data.

One important conclusion that can be reached by the analysis of the case study results is the fact that the explicit consideration of late assistance arrivals can indeed have an impact on the emergency vehicles' location decisions. It is possible to improve coverage by relocating the existing vehicles in the geographical area under study, guaranteeing a more equitable access to this essential service, without increasing the total number of vehicles. Having a model that is capable of better representing the real situation can be an added value for decision-making support.

The heuristic approach presented turned out to be able to calculate a high-quality solution, with a very small gap. This heuristic approach has the advantage of being well understood by decision-makers, since it relies on the optimal decisions calculated for small instances corresponding to different scenarios, and it is fairly easy to make an analogy with the situation where a set of experts is called to give their opinion. These experts can have different opinions on some aspects, share the same points of view in other aspects, and the final decision will consider what most of the experts agree with. The presented iterative and incremental way of building the solution has also the advantage of having a computational time that grows linearly with the number of scenarios considered.

4.3. Future work

In this model there are some situations that are not being taken into account. One example is the possibility of vehicles changing their locations during the day, or to consider different locations on different days of the week, or periods of the year. It could also be interesting to introduce the possibility of locating vehicles outside existing bases, if that contributes to increased coverage. Changes in the location of vehicles are very difficult to implement due to the disagreement of the professionals involved, and the resistance to these dynamic changes would even be stronger. Changing the location of a vehicle implies other logistic changes, namely considering the replacement of the materials consumed in the assistance activity, which are possible to address, but difficult to organize. Nevertheless, it can be interesting to understand if these possibilities could contribute or not to the improvement of the emergency episodes coverage.

The developed work motivates further research challenges to be pursuit. The difficulty experienced by the general solver justified the use of the heuristic approach described. Other possibilities could be considered, like using decomposition methods. This type of problems has clearly an inherently multiobjective nature, that was not explicitly tackled in this work.

CRediT authorship contribution statement

José Nelas: Conceptualization, Methodology, Investigation, Writing – original draft. **Joana Dias:** Methodology, Writing – review & editing, Supervision.

Declaration of interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This study has been funded by national funds, through FCT, the Portuguese Science Foundation, under project UIDB/00308/2020 and with the collaboration of Coimbra paediatric Hospital – Coimbra Hospital and University Centre, and INEM.

Appendix

Acronymous
AA: Assistance Ambulance
BU: Basic Urgency
CH: Central Hospital
GH: General Hospital
ILSA: Immediate Life Support Ambulance
INEM: Instituto Nacional de Emergência Médica
MEA: Medical Emergency Ambulance
MEM: Medical Emergency Motorcycle
MERV: Medical Emergency and Resuscitation Vehicles
PF: Professional Firefighters
RB: Regional Base
RC: Red Cross
RH: Regional Hospital
VF: Volunteer firefighters

References

[1] National Association of State EMS Officials, “National EMS scope of practice model 2019,” 2019.
[2] INEM. Manual de suporte avançado de vida. Lisboa: INEM; 2019.
[3] Chanta S, Mayorga ME, McLay LA. Improving emergency service in rural areas: a bi-objective covering location model for EMS systems. Ann Oper Res 2014;221(1): 133–59.
[4] Maxwell MS, Restrepo M, Henderson SG, Topaloglu H. Approximate dynamic programming for ambulance redeployment. INFORMS J Comput 2010;22(2): 266–81.
[5] Fitch J. Response times: myths, measurement & management. JEMS 2005;30(9): 47—56.
[6] Macharia WM, Njeru EK, Muli-Musiime F, Nantulya V. Severe road traffic injuries in Kenya, quality of care and access. Afr Health Sci 2009;9(2):118–24.
[7] Mahama MN, Kenu E, Bandoh DA, Zakariah AN. Emergency response time and pre-hospital trauma survival rate of the national ambulance service, Greater Accra (January - December 2014). BMC Emerg Med 2018;18(1):3–10.
[8] Sánchez-Mangas R, García-Ferrrer A, De Juan A, Arroyo AM. The probability of death in road traffic accidents. How important is a quick medical response? Accid Anal Prev 2010;42(4):1048–56.
[9] Ma L, Zhang H, Yan X, Wang J, Song Z, Xiong H. Smooth associations between the emergency medical services response time and the risk of death in road tra ffi c crashes. J Transp Heal 2019;12:379–91. August 2018.
[10] Hess EP, White RD. Optimizing survival from out-of-hospital cardiac arrest: clinical review. J Cardiovasc Electrophysiol 2010;21(5):590–5.
[11] Nadarajan GD, et al. Global resuscitation alliance utstein recommendations for developing emergency care systems to improve cardiac arrest survival. Resuscitation 2018;132:85–9. August.
[12] Sumrit and Thongsiriruengchai. An optimization model for advanced life support ambulance facility location problem. Proceedings 2020;39(1):10.
[13] Yang W, Su Q, Zhou M, Qin X. Ambulance allocation considering the spatial randomness of demand. Comput Ind Eng 2020;139:106202. November 2019.
[14] Park H, Shafahi A, Haghani A. Considering secondary incidents on freeways. IEEE Trans Intell Transp Syst 2016;17(9):2528–40.
[15] INEM, “Indicadores de Desempenho,” 2021.
[16] Ahmadi-Javid A, Seyedi P, Syam SS. A survey of healthcare facility location. Comput Oper Res 2017;79:223–63.
[17] Bélanger V, Ruiz A, Soriano P. Recent optimization models and trends in location, relocation, and dispatching of emergency medical vehicles. Eur J Oper Res 2019; 272(1):1–23.
[18] Boloori Arabani A, Farahani RZ. Facility location dynamics: an overview of classifications and applications. Comput Ind Eng 2012;62(1):408–20.
[19] Farahani RZ, Asgari N, Heidari N, Hosseininia M, Goh M. Covering problems in facility location: a review. Comput Ind Eng 2012;62(1):368–407.
[20] Andrade LACG, Cunha CB. An ABC heuristic for optimizing moveable ambulance station location and vehicle repositioning for the city of So Paulo. Int Trans Oper Res 2015;22(3):473–501.
[21] Bertsimas D, Ng Y. Robust and stochastic formulations for ambulance deployment and dispatch. Eur J Oper Res 2019;279(2):557–71.
[22] Nelas J, Dias J. Optimal Emergency Vehicles Location: an approach considering the hierarchy and substitutability of resources. Eur J Oper Res 2020. to appear.
[23] Peng C, Delage E, Li J. Probabilistic envelope constrained multiperiod stochastic emergency medical services location model and decomposition scheme. Transp Sci 2020;54(6):1471–94.
[24] Yoon S, Albert LA, White VM. A stochastic programming approach for locating and dispatching two types of ambulances. Transp Sci 2021;55(2):275–96.
[25] Zhang R, Zeng B. Ambulance deployment with relocation through robust optimization. IEEE Trans Autom Sci Eng 2019;16(1):138–47.
[26] Mohri SS, Haghshenas H. An ambulance location problem for covering inherently rare and random road crashes. Comput Ind Eng 2021;151:106937. October 2020.
[27] Boujemaa R, Jebali A, Hammami S, Ruiz A. Multi-period stochastic programming models for two-tiered emergency medical service system. Comput Oper Res 2020; 123:104974.

[28] Akıncılar A, Akıncılar E. A new idea for ambulance location problem in an environment under uncertainty in both path and average speed: absolutely robust planning. Comput Ind Eng 2019;137:106053. July 2018.
[29] Wajid S, Nezamuddin N, Unnikrishnan A. Optimizing ambulance locations for coverage enhancement of accident sites in South Delhi. Transp Res Procedia 2020; 48:280–9.
[30] Bélanger V, Lanzarone E, Nicoletta V, Ruiz A, Soriano P. A recursive simulation-optimization framework for the ambulance location and dispatching problem. Eur J Oper Res 2020;286(2):713–25.