



HAL
open science

Inhomogeneous Poisson Process for Ambulance Relocation

Félicien Hêche, Oussama Barakat, Thibaut Desmettre, Tania Marx, Stephan Robert-Nicoud

► **To cite this version:**

Félicien Hêche, Oussama Barakat, Thibaut Desmettre, Tania Marx, Stephan Robert-Nicoud. Inhomogeneous Poisson Process for Ambulance Relocation. IWCTS '23: 16th ACM SIGSPATIAL International Workshop on Computational Transportation Science, Nov 2023, Hamburg, Germany. pp.1-8, 10.1145/3615895.3628172 . hal-04661768

HAL Id: hal-04661768

<https://hal.science/hal-04661768v1>

Submitted on 25 Jul 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License



Inhomogeneous Poisson Process for Ambulance Relocation

Félicien Hêche
felicien.heche@hes-so.ch
School of Engineering and
Management, University of Applied
Sciences and Arts Western
Switzerland (HES-SO)
Yverdon-les-Bains, Switzerland

Oussama Barakat
oussama.barakat@univ-fcomte.fr
Nanomedecine Lab, Imagery and
Therapeutics, University of
Franche-Comté
Besançon, France

Thibaut Desmettre
Thibaut.Desmettre@hcuge.ch
Emergency Department, Hôpitaux
Universitaires de Genève (HUG)
Genève, Switzerland

Tania Marx
tmarx@chu-besancon.fr
Emergency Department, Centre
Hospitalier Universitaire (CHU)
Besançon, France

Stephan Robert-Nicoud
stephan.robert@hes-so.ch
School of Engineering and
Management, University of Applied
Sciences and Arts Western
Switzerland (HES-SO)
Yverdon-les-Bains, Switzerland

ABSTRACT

This paper proposes a statistical approach for real-time ambulance relocation in an Emergency Medical Service (EMS) system. First, based on the idea of measuring the probability that there is no ambulance available for a life-threatening call at the closest ambulance station, the risk of the environment is introduced. Then, a method for ambulance relocation to minimize this risk is developed. To assess the proposed approach, historical data spanning six years provided by the Centre Hospitalier Universitaire Vaudois (CHUV) are used. Data from 2015 to 2020 is utilized to build our method and the data from 2021 is used for the evaluation. The use of real-world data and an API to estimate travel times makes our experiments representative of real-world situations. In all experiments, our approach enables us to significantly decrease the risk of the environment. Furthermore, the proposed method reduces the mean response time by up to 30 seconds. Finally, the computation time is negligible and through the adjustment of specific hyperparameters, it becomes feasible to control the frequency of resource relocations throughout the day. This adjustment offers the flexibility to tailor the strategy according to the specific requirements of each EMS system. Moreover, our method demonstrates particular utility in situations with limited resources, which makes it especially valuable in crisis situations. These observations collectively suggest that our method has the potential to enhance the quality of EMS and, consequently, save lives.

CCS CONCEPTS

• **Computing methodologies** → **Model development and analysis.**



This work is licensed under a Creative Commons Attribution International 4.0 License.

IWCTS '23, November 13, 2023, Hamburg, Germany

© 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0357-7/23/11.

<https://doi.org/10.1145/3615895.3628172>

KEYWORDS

Real-time ambulance relocation, Inhomogeneous Poisson process, Emergency Medical Service (EMS)

ACM Reference Format:

Félicien Hêche, Oussama Barakat, Thibaut Desmettre, Tania Marx, and Stephan Robert-Nicoud. 2023. Inhomogeneous Poisson Process for Ambulance Relocation. In *16th ACM SIGSPATIAL International Workshop on Computational Transportation Science (IWCTS '23)*, November 13, 2023, Hamburg, Germany. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3615895.3628172>

1 INTRODUCTION

Emergency Medical Services (EMS) are essential medical services that provide urgent pre-hospital treatment and transportation for injured individuals to intensive care. EMS utilizes a finite number of ambulances which are stationed at various locations across the region of interest when not in use. Ambulance deployment is carried out at two different levels: at the strategic level, where the location of ambulance stations is determined, and at the operational level, where ambulances are dispatched and relocated. It is common to classify emergency calls into three priority levels: Priority 1 for the most urgent life-threatening calls, Priority 2 for urgent but not life-threatening calls, and Priority 3 for non-urgent calls. Readers interested in gaining a deeper understanding of how these priority levels are established in practice are invited to consult Dami et al. [6]. In this paper, emergency calls will refer to calls of Priority 1 and 2. Priority 1 emergencies will also be referred to as vital emergencies. Relocating ambulances through the stations could be useful to decrease the response time (i.e. the elapsed time from EMS dispatch notification to the time of arrival on scene), which might have a significant impact on the survival rate [9], [29].

This paper presents a statistical method for real-time ambulance relocation. This approach is simple and based on statistical considerations. As illustrated in Figure 1, 2, and 3, the number of emergency calls varies significantly based on the time of the day and the location but not so much based on other parameters, such as the day of the week or the month. Based on these observations, we choose to associate one stochastic process with each station to model the








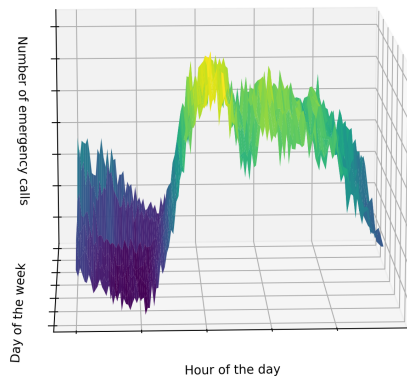


Figure 1: Number of emergency calls for a time interval of ten minutes for the region covered by the CHUV dispatch center for the year 2016-2021. The x-axis represents the time and the y-axis the day of the week.

Figure 2: Number of emergency calls for a time interval of ten minutes for the region covered by the CHUV dispatch center for the year 2016-2021. The x-axis represents the time and the y-axis the month.

number of vital emergency calls. Specifically, and following common practice [31], [26], we use inhomogeneous Poisson processes. By definition, inhomogeneous Poisson processes are stochastic processes that satisfy two assumptions. Roughly speaking, the first one supposes that the probability of receiving one call within a time interval can be approximated using a time-dependent event rate. The second one assumes that two calls can not occur at the same time. These assumptions will be formally presented in Section 3.1. The inherent suitability of these assumptions for our problem justifies the use of inhomogeneous Poisson processes.

Combining these processes and the idea of measuring the probability that there is no ambulance available for a life-threatening call at the nearest ambulance station, we introduce the *risk* of the environment and we develop a method that relocates resources to minimize this risk. Our approach avoids the use of Mixed Integer Linear Programming (MILP) techniques, known for their high computational demands, making the proposed method computationally efficient. We evaluate the performance of our approach

using historical data provided by the Centre Hospitalier Universitaire Vaudois (CHUV), Switzerland. In our experiments, we chose to use an API to estimate the travelling time. The time spent on the site of each incident by the medical staff has been estimated using the historical dataset. These choices have been made to ensure that our experiments are representative of real-world situations. The results suggest that the proposed method is able to reduce the risk of the environment and the expected response time. Through the adjustment of specific hyperparameters, it becomes feasible to control the number of resource relocations throughout the day. This adjustment offers the flexibility to tailor the strategy according to the specific requirements of each EMS system. Moreover, our method appears to be particularly useful in situations with limited resources, which makes it especially valuable in crisis situations. All these observations suggest that our method could increase the quality of EMS and thus, potentially save lives.

This paper is organized as follows. In the next section, we provide a literature review. Then, in Section 3, we define our problem and present our method formally. In Section 4, we evaluate our approach




Figure 3: Locations of emergency calls of for the years 2016 to 2021

using historical data. Finally, the conclusion of our research is given in Section 5.

2 LITERATURE REVIEW

Given the critical role that EMS plays in the medical system, it is not surprising that researchers have dedicated significant attention to studying efficient resource [25], [11], [27], [3], [18], [2], [23]. In particular, ambulance management has raised a lot of attention.

At the strategic level, it is common to use a Mixed Integer Linear Program (MILP) to determine the optimal location of ambulance stations. These methods can be separated into two different families: the deterministic and the probabilistic methods. In the deterministic approach, authors do not take into account some stochastic considerations and assume that ambulances are always available for dispatch. For example, the Location Set Covering Problem [28] minimizes the number of stations that are necessary to cover all demand points within a certain response time. The Maximal Covering Locations Problem [5] maximizes the total demand covered within a certain response time or travel distance given a fixed number of stations. The Backup Coverage Problem [16] maximizes the demand that is covered twice within a certain response time or travelling distance. In Gendreau et al. [12], a heuristic tabu search is introduced to solve the Double Standard Model, which considers two radii $r_1 < r_2$ and aims to maximize the demand covered by at least two ambulances within r_1 while ensuring that all demand is covered within r_2 . This model is extended in Doerner et al. [8] by introducing a limit on the number of inhabitants served per ambulance and a heuristic is introduced to solve it.

With the probabilistic models, the possibility that an ambulance is busy is taken into account. Some of these models aim to determine the optimal ambulance location but also to compute the number of ambulances assigned to each station. For example, the Maximum Expected Covering Location Problem (MEXCLP) [7] attempts to

maximize an expected coverage that takes into account the average fraction of time that an ambulance is unavailable (i.e. the busy fraction of the ambulances). The Maximum Availability Location Problem [24] maximizes the demand covered by ambulances with a given probability. The busy fraction is assumed to be the same for all potential stations. In Erkut et al. [10], the probability of survival is modeled, and then, the maximal survival location problem is introduced. Finally, Chuang and Lin [4] introduces a model to solve a double standard coverage ambulance location under a probabilistic situation.

At the operational level, only ambulance dispatch and relocation can be considered. As mentioned in the introduction, it is common to classify emergency calls into three priority levels. Considering these priority levels, several dispatch strategies have been developed. For instance, Bandara et al. [1] proposes a heuristic algorithm to dispatch ambulances to increase the patient's survival probability. Nasrollahzadeh et al. [22] formulates the real-time ambulance dispatching and relocation problem as a stochastic dynamic problem. An approximate dynamic programming problem is developed to solve this problem. Another study [30] introduces the Testing Interface For Ambulance Research (TIFAR) simulation tool that can be used to evaluate the effectiveness of different dispatch strategies.

Relocating ambulances within a specific geographic region can enhance the coverage and response time of EMS. However, in this context, computation time has to be reasonably low, requiring the adoption of distinct methods from those used at the strategic level. To overcome this challenge, Gendreau et al. [13] uses a tabu search algorithm for real-time ambulance relocation. In Gendreau et al. [14], a MILP is introduced for locating and relocating physicians' cars in the Montreal area. The small number of such vehicles (between three and six) makes this approach feasible. Another method [20] employs a dynamic model for the ambulance relocation problem. To reduce the problem's dimension and computation time, approximations that depend on specific parameters calibrated using simulation, are used. Finally, Naoum-Sawaya and Elhedhli [21] introduces a two-stage stochastic program for real-time ambulance relocation.

Inhomogeneous Poisson processes are stochastic processes that have been used to model various situations. For instance, they have been applied for call center management [31], [26], for software reliability [17], for automobile warranty claim [19] but also for predicting the short-term public transport demand [15].

3 STATISTICAL APPROACH

This section primarily focuses on the formal definition of the risk. Next, we present a method for minimizing this risk through ambulance relocation. Finally, we provide details on the practical implementation of this approach.

3.1 Risk

As the response time has a bigger influence on the survival rate for vital emergencies, we mainly focus on these kinds of incidents. First, let S denote the set of all stations. For a given station $a \in S$, the region covered by a corresponds to villages whose closest station, in terms of travelling time, is a . $N_{t,a}$ denotes the number of vital

incidents in the time interval $(0, t]$ in the region covered by the station a . Furthermore, given $a, b \in S$, we use $d(a, b)$ to represent the travelling time between these two stations. Finally, we use $m_{t,a}$ to denote the number of available ambulances in station a at time t .

Using these notations, the following assumptions for all stations a and all time t are made

- (1) $\mathbb{P}[N_{t+\epsilon,a} - N_{t,a} = 1] = \lambda_a(t)\epsilon + o(\epsilon)$
- (2) $\mathbb{P}[N_{t+\epsilon,a} - N_{t,a} \geq 2] = o(\epsilon)$

These statements imply that vital emergencies form an inhomogeneous Poisson process with a density function $\lambda_a(t)$. In other words, $N_{t_M,a} - N_{t_0,a}$ is a Poisson process of parameter $\Lambda(t_0, t_M, a)$ where

$$\Lambda(t_0, t_M, a) = \int_{t_0}^{t_M} \lambda_a(t) dt$$

Then, for a given time t and station a , we introduce a decomposition $t = t_0 < t_1 < \dots < t + \gamma(a) = t_{T_a}$ where $\gamma(a) = \min_{b \in S \setminus \{a\}} d(a, b)$, and each timestep t_i corresponds to time where a previously engaged ambulance becomes available again. After that, we introduce a random variable $\mathcal{R}_{t,a}$ defined as

$$\mathcal{R}_{t,a} = \sum_{j=0}^{t_{T_a}-1} \mathbb{1}_{\{N_{t_{j+1}} - N_{t_j} > m_{t_j,a}\}}$$

In other words, if $\mathcal{R}_{t,a} \geq 1$, we are in a situation where there will be more vital incidents than available ambulances in station a , and we will not have enough time to relocate ambulance to this station. Obviously, this is the kind of situation that needs to be avoided. Finally, we define the *risk* of the environment at time t as $\mathbb{E}[\mathcal{R}_t]$, where $\mathcal{R}_t = \sum_{a \in S} \mathcal{R}_{t,a}$.

3.2 Ambulances relocation

Now, we present a method for ambulance relocation. The main idea of our approach is straightforward: we relocate resources if it decreases the risk more than a given threshold.

Let us formalize this idea. First, we denote $\mathcal{R}_{t_{ab}}(a, b)$ the risk in time $t_{ab} := t + d(a, b)$ assuming we relocate an ambulance from station b to station a in time t . Then, observe that

$$\begin{aligned} & \mathbb{E}[\mathcal{R}_{t_{ab}}] - \mathbb{E}[\mathcal{R}_{t_{ab}}(a, b)] \\ &= \mathbb{E}\left[\sum_{c \in S \setminus \{a, b\}} \mathcal{R}_{t_{ab},c}\right] + \mathbb{E}[\mathcal{R}_{t_{ab},a}] + \mathbb{E}[\mathcal{R}_{t_{ab},b}] \\ & \quad - \mathbb{E}\left[\sum_{c \in S \setminus \{a, b\}} \mathcal{R}_{t_{ab},c}\right] - \mathbb{E}[\mathcal{R}_{t_{ab},a}(a, b)] - \mathbb{E}[\mathcal{R}_{t_{ab},b}(a, b)] \\ &= \mathbb{E}[\mathcal{R}_{t_{ab},a}] - \mathbb{E}[\mathcal{R}_{t_{ab},a}(a, b)] + \mathbb{E}[\mathcal{R}_{t_{ab},b}] - \mathbb{E}[\mathcal{R}_{t_{ab},b}(a, b)] \\ &= \sum_{j=0}^{t_{T_a}-1} \mathbb{E}[\mathbb{1}_{\{N_{t_{j+1}} - N_{t_j} > m_{t_j,a}\}}] - \sum_{j=0}^{t_{T_a}-1} \mathbb{E}[\mathbb{1}_{\{N_{t_{j+1}} - N_{t_j} > m_{t_j,a} + 1\}}] \\ & \quad + \sum_{i=0}^{t_{T_b}-1} \mathbb{E}[\mathbb{1}_{\{N_{t_{i+1}} - N_{t_i} > m_{t_i,b}\}}] - \sum_{i=0}^{t_{T_b}-1} \mathbb{E}[\mathbb{1}_{\{N_{t_{i+1}} - N_{t_i} > m_{t_i,b} - 1\}}] \end{aligned}$$

$$\begin{aligned} &= \sum_{j=0}^{t_{T_a}-1} \mathbb{P}[N_{t_{j+1},a} - N_{t_j} > m_{t_j,a}] \\ & \quad - \sum_{j=0}^{t_{T_a}-1} \mathbb{P}[N_{t_{j+1},a} - N_{t_j} > m_{t_j,a} + 1] \\ & \quad + \sum_{i=0}^{t_{T_b}-1} \mathbb{P}[N_{t_{i+1},b} - N_{t_i} > m_{t_i,b}] \\ & \quad - \sum_{i=0}^{t_{T_b}-1} \mathbb{P}[N_{t_{i+1},b} - N_{t_i} > m_{t_i,b} - 1] \\ &= \sum_{j=0}^{t_{T_a}-1} \mathbb{P}[N_{t_{j+1},a} - N_{t_j} = m_{t_j,a} + 1] \\ & \quad - \sum_{i=0}^{t_{T_b}-1} \mathbb{P}[N_{t_{i+1},b} - N_{t_i,b} = m_{t_i,b}] \end{aligned}$$

Moreover, for numerical considerations, we use the log probability instead of the probability. Therefore, we introduce a function h defined as

$$\begin{aligned} h(a, b, t) := & \sum_{j=0}^{T_a-1} \log\left(\mathbb{P}[N_{t_{j+1},a} - N_{t_j,a} = m_{t_j,a} + 1]\right) \\ & - \sum_{i=0}^{T_b-1} \log\left(\mathbb{P}[N_{t_{i+1},b} - N_{t_i,b} = m_{t_i,b}]\right) - \beta(a, b) \end{aligned} \quad (1)$$

where $\beta(a, b)$ is a hyperparameter that models the cost associated with relocating an ambulance from station b to station a . In other words, $\beta(a, b)$ is a value that determines how easily we allow ambulance relocation from b to a . This parameter should be carefully tuned according to the strategic choices of the EMS.

Finally, at each time t we suggest relocating ambulances from station b to a if $h(a, b, t) > 0$. If there exist multiple such pairs, we propose considering only the one that maximizes h .

3.3 Implementation details

Although the ambulance relocation method presented above is quite simple, there are some implementation details that need to be discussed.

First, to compute the region covered by each station, we estimate the closest station of each village using the Openrouteservice API¹. If this API is unable to estimate this travelling time (due to errors in the GPS coordinates for example), we use the Euclidean distance.

Then, we need to be able to compute probabilities of the form $\mathbb{P}[N_{t_{j+1},a} - N_{t_j,a} = m_{t_j,a}]$. Recall that $N_{t_{j+1},a} - N_{t_j,a}$ is a Poisson process with parameter $\Lambda(t_j, t_{j+1}, a)$. Therefore, to compute these probabilities, we need to estimate this parameter. With this aim, we separate our time interval into subintervals $t_j := t^0 < \dots < t^k < \dots < t^N := t_{j+1}$ where $t^k = t^0 + k\epsilon$ and ϵ is the smallest considered timestep. For example, in our implementation, we set ϵ to one minute. Then, since $\Lambda(t_j, t_{j+1}, a) = \sum_{k=0}^{N-1} \Lambda(t^k, t^{k+1}, a)$, we only need to estimate the parameters of the random variable

¹<https://openrouteservice.org/>

$N_{t_{j+1},a} - N_{t_j,a}$. To achieve this objective, we use the maximum likelihood estimator, which can be computed using the historical dataset. Recall that the maximum likelihood estimator of a Poisson process is the empirical mean. However, to avoid taking the logarithm of zeros values, an artificial incident for each time step and each station is added. With these considerations, we are able to estimate probabilities of the form $\mathbb{P}[N_{t_{j+1},a} - N_{t_j,a} = m_{t_j,a}]$.

To compute the return time of an ambulance dispatched in time t , we adopt the following idea. First, the travelling time between the station a and the emergency location s , denoted $d(a, s)$, is estimated using the Openrouteservice API as above. If this API is not able to estimate this travelling time, we consider a default value of 20 minutes. Finally, the estimated return time is $t + d(a, s) + \bar{t}$ where \bar{t} represents the mean of the time spent on the site of the incident.

Lastly, we need to choose the value of the hyperparameter β . We suggest considering $\beta(a, b) = \gamma d(a, b)$, with $\gamma \in \mathbb{R}$. A small γ reduces the risk but increases the number of ambulance relocations. As we aim to minimize the risk of the environment while avoiding making too many ambulance relocations, finding an appropriate γ may not be straightforward. Therefore in our experiments, we test different values of γ .

4 EXPERIMENTS

In this section, we measure the performance of the proposed method in different real-world scenarios. First, we present our evaluation process. Then, our results are presented and discussed.

4.1 Experimental setup

To build and evaluate our models, historical data provided by the CHUV dispatch center are used. This dispatch center is responsible for managing ambulances in the canton of Vaud, in the canton of Neuchâtel, as well as some small parts of the canton of Fribourg and Genève. In total, about one million people are living in this




Figure 4: Region covered by the CHUV dispatch center. The canton of Vaud is in red, the canton of Neuchâtel is in green and the yellow portion represents the area of the canton of Fribourg covered by this center.

area. There are 18 ambulance stations located throughout it. The region covered by this dispatch center is shown in Figure 4 and the station locations can be observed in Figure 5.

Different files regarding incidents and resources history from the year 2014 to 2021 were provided. For each incident, some information such as time and address is logged. Ambulances can take several

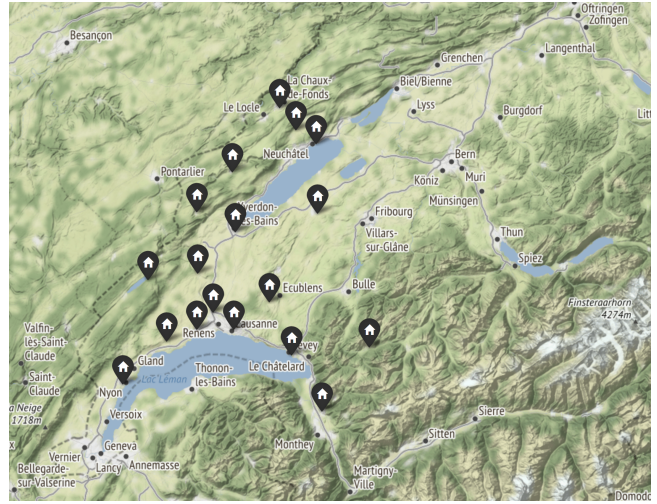


Figure 5: Ambulance locations of the CHUV dispatch center.

statuses, like for instance 'Commitment request', 'Out', 'On site'. Whenever these statuses change, relevant information is saved. We use historical data from the year 2015 to 2020 to build our method and the data from the year 2021 for the evaluations. To estimate the mean time spent on the site of the incident \bar{t} , data from the year 2015 until the end of 2020 has been used. Specifically, for each incident, we estimated the dispatched ambulance using the GPS coordinate contained in the historical dataset. Then, the estimated time spent on the incident site corresponds to the time interval between the status 'End' and 'On site'. We were able to estimate this time interval for 101'060 incidents.

Using these files and estimation, we conduct several experiments, denoted Historical and Distribution_n.

Historical. For each day of the year 2021, we generated incidents based on the historical dataset provided by the CHUV dispatch center. Therefore, the time and incident locations used during our evaluation process correspond to real incidents that occurred during the year 2021. Ambulances are allocated to emergencies using a greedy method, i.e. the closest available ambulance is dispatched. The time spent on the site of the incident is estimated using the historical dataset. If we are unable to estimate it, we consider the mean value \bar{t} . Ambulances are relocated following the method introduced in Section 3.2. At the beginning of each day, we assign to each station the number of ambulances dispatched from the corresponding station in our historical dataset and in the corresponding day.

Distribution_n. Note that the experiment presented above suffers from some limitations regarding the number of ambulances. Indeed, if an ambulance was available but never dispatched, we would not take it into account. Additionally, it is possible that some ambulances are only available at a certain time of the day. To address these limitations and test our method in other settings, we introduce an evaluation called Distribution_n where $n \in \mathbb{N}$. The difference with the experiment presented above concerns the number and locations of ambulances available at the beginning of each day. More

Table 1: Results of methods for managing strategic moves for the Historical experiment. We bold the method with the smallest mean response time and the smallest mean cumulative risk.

| γ | Risk | Relocation | Response time |
|-----------|---------------|---------------|-----------------|
| $+\infty$ | 2.99 | 0.0 | 09:44.46 |
| 0 | 0.1335 | 113.386 | 09:19.86 |
| 0.002 | 0.1055 | 70.145 | 09:15.96 |
| 0.004 | 0.2840 | 37.381 | 09:16.19 |
| 0.006 | 0.5931 | 23.573 | 09:16.99 |
| 0.008 | 0.9215 | 17.381 | 09:17.66 |
| 0.010 | 1.1977 | 13.385 | 09:19.64 |
| 0.012 | 1.3546 | 10.523 | 09:21.19 |
| 0.014 | 1.5658 | 8.899 | 09:23.58 |
| 0.016 | 1.7402 | 7.474 | 09:25.39 |
| 0.018 | 1.9332 | 6.540 | 09:28.28 |
| 0.02 | 2.1481 | 5.616 | 09:31.01 |

precisely, we denote \mathcal{M} the set of all incident locations which are contained in our dataset. Then, for each station a , we introduce a subset $\mathcal{M}_a := \{m \in \mathcal{M} \mid \|m - a\|_2 < \|m - v\|_2 \forall v \in S, v \neq a\}$ and we associate a distribution $p_a := \frac{|\mathcal{M}_a|}{|\mathcal{M}|}$. Finally, at the beginning of each day, we assign $\lfloor p_a n \rfloor$ ambulances to each station a .

Experiments were performed on a laptop computer (Intel Core i7-10750H, 32GB RAM) using open-source software. We used Python (version 3.8) with libraries such as NumPy and Pandas. Performing one experiment takes about 4.5 hours and the computation of one potential resource relocation (i.e. the computation of $h(a, b, t)$ for each $a, b \in S$) takes about 0.03 seconds.

4.2 Results

As previously discussed, we consider β of the form $\beta(a, b) = \gamma d(a, b)$. For each experiment, we test several γ values. Note that setting $\gamma = +\infty$ implies that no ambulance relocation will be conducted throughout the day. To measure the performance of our approach we use the cumulative risk $\sum_t \mathbb{E}[\mathcal{R}_t]$ where the sum is taken after each dispatch, the number of ambulance relocation and the response time. We compute these values for each day, and then, we take the mean. Note that the cumulative risk computed here is not the same as the risk estimation presented in Section 3.2. Indeed, previously, the time spent on the site of an incident was approximated using a constant, but in our evaluations, it is derived from the historical dataset. Therefore, it should be more representative of a real-world situation. Complete results for the Historical experiment can be found in Table 1, while Tables 2, 3, 4, present results regarding Distributional₄₂, Distributional₅₂, Distributional₆₂ respectively.

4.3 Discussion

In this section, we discuss the results presented above.

Firstly, we would like to highlight that our method consistently decreases the cumulative risk, regardless of the value of γ or the experiment. Moreover, and this may be more surprising, it also

Table 2: Results of methods for managing strategic moves on Distribution₄₂ benchmark. We bold the method with the smallest mean response time and the smallest mean cumulative risk.

| γ | Risk | Relocation | Response time |
|-----------|---------------|---------------|-----------------|
| $+\infty$ | 3.4957 | 0.0 | 10:07.57 |
| 0 | 1.2020 | 113.386 | 10:05.64 |
| 0.002 | 0.8115 | 55.644 | 09:37.26 |
| 0.004 | 2.0647 | 30.039 | 09:42.65 |
| 0.006 | 2.5534 | 15.794 | 09:53.06 |
| 0.008 | 2.7021 | 10.488 | 09:56.03 |
| 0.010 | 2.7852 | 8.159 | 09:56.86 |
| 0.012 | 2.9301 | 6.074 | 09:57.92 |
| 0.014 | 2.9935 | 4.852 | 09:58.25 |
| 0.016 | 3.1002 | 4.140 | 10:00.25 |
| 0.018 | 3.1791 | 3.179 | 10:01.95 |
| 0.02 | 3.2408 | 3.287 | 10:03.56 |

Table 3: Results of methods for managing strategic moves on Distribution₅₂ benchmark. We bold the method with the smallest mean response time and the smallest mean cumulative risk.

| γ | Risk | Relocation | Response time |
|-----------|---------------|---------------|-----------------|
| $+\infty$ | 2.2832 | 0.0 | 09:38.58 |
| 0 | 0.2647 | 113.386 | 09:27.11 |
| 0.002 | 0.2101 | 64.841 | 09:18.50 |
| 0.004 | 1.1496 | 32.323 | 09:25.18 |
| 0.006 | 1.7413 | 20.186 | 09:30.21 |
| 0.008 | 1.9184 | 13.649 | 09:33.45 |
| 0.010 | 1.9954 | 10.115 | 09:35.19 |
| 0.012 | 2.0173 | 7.865 | 09:35.53 |
| 0.014 | 2.0566 | 6.479 | 09:36.04 |
| 0.016 | 2.1557 | 5.564 | 09:37.00 |
| 0.018 | 2.1965 | 4.824 | 09:37.27 |
| 0.02 | 2.2178 | 4.315 | 09:37.46 |

reduces the response time. For example, for Distribution₄₂ and $\gamma = 0.002$, our method can decrease the mean response time by up to 30 seconds. Moreover, there is an important correlation between the cumulative risk and the response time. For instance, for the Historical experiment, we have a correlation of 0.91 between these two variables. This observation suggests that reducing the cumulative risk can effectively decrease the response time in practical scenarios. Then, we notice that gaps between our method and the case where no ambulance relocation occurs are more significant in situations with limited resources. This finding suggests that our approach is particularly useful in crisis situations where many ambulances are required, such as a large traffic incident.

Secondly, it is worth noting that for all experiments the most effective ambulance relocation method in terms of response time and cumulative risk is achieved with $\gamma = 0.002$. Interestingly, the

Table 4: Results of methods for managing strategic moves on *Distribution*₆₂ benchmark. We bold the method with the smallest mean response time and the smallest mean cumulative risk.

| γ | Risk | Relocation | Response time |
|-----------|---------------|---------------|-----------------|
| $+\infty$ | 0.5885 | 0.0 | 09:22.10 |
| 0 | 0.0789 | 113.386 | 09:19.17 |
| 0.002 | 0.0373 | 67.761 | 09:15.61 |
| 0.004 | 0.1472 | 36.326 | 09:15.73 |
| 0.006 | 0.2112 | 23.063 | 09:16.81 |
| 0.008 | 0.3405 | 15.909 | 09:18.40 |
| 0.010 | 0.4386 | 12.931 | 09:19.37 |
| 0.012 | 0.4630 | 10.449 | 09:19.69 |
| 0.014 | 0.4763 | 9.495 | 09:19.82 |
| 0.016 | 0.4843 | 8.600 | 09:19.91 |
| 0.018 | 0.4928 | 7.816 | 09:19.94 |
| 0.020 | 0.5047 | 7.035 | 09:20.31 |

cumulative risk is smaller for $\gamma = 0.002$ than $\gamma = 0.0$. While this observation may first seem surprising, it can be easily explained. Indeed, with $\gamma = 0$, resources are relocated if it will reduce the estimated risk, even if the gain is small. However, the estimated risk is not an exact match for the risk calculated in our evaluations. Recall that, in our estimations, the time spent at the scene of an incident is approximated using a constant, whereas, in our evaluation, it is based on historical data. Therefore, small estimated gains can not be trusted, and thus, making resource relocation based on these estimations could actually be counter-productive.

Finally, it is interesting to note that the observations discussed above are consistent across all experiments, indicating that our method is robust and, thus, could be successfully applied in various EMS systems.

5 CONCLUSION

In this paper, we present a novel method for real-time ambulance relocation. First, using an inhomogeneous Poisson process, we model Priority 1 calls, and then, we introduce the notion of risk of the environment. Based on this definition, we develop a method for resource relocation that minimizes this risk. By adjusting the β parameter, it is possible to manage the number of resource relocations throughout the day and avoid certain specific resource relocations.

To evaluate our method, we use a historical dataset provided by the CHUV dispatch center. Years from 2015 to 2020 are used to build our method and the year 2021 is used for the evaluation. In our experiment, generated incidents are based on historical emergencies with corresponding time and location, and travelling time is computed using the Openrouteservice API.

In all experiments, our method significantly reduces the expected risk. Furthermore, our results suggest that minimizing the proposed risk significantly decreases the response time. Moreover, our method appears to be particularly useful in situations with limited resources, which makes it especially valuable in crisis situations.

All these observations suggest that our method could increase the quality of EMS and thus potentially save lives.

ACKNOWLEDGMENTS

We would like to thank the CHUV dispatch center and more specifically Dr. Fabrice Dami for providing historical data used in this study.

REFERENCES

- [1] Damitha Bandara, Maria E Mayorga, and Laura A McLay. 2014. Priority dispatching strategies for EMS systems. *Journal of the Operational Research Society* 65, 4 (2014), 572–587. <https://doi.org/10.1057/jors.2013.95>
- [2] JJM Black, ME Ward, and DJ Lockey. 2004. Appropriate use of helicopters to transport trauma patients from incident scene to hospital in the United Kingdom: an algorithm. *Emergency Medicine Journal* 21, 3 (2004), 355–361. <https://doi.org/10.1136/emj.2002.004473>
- [3] Maurizio Brughieri, Cesare Cardani, and Matteo Putzu. 2016. Optimizing the Location of Helicopter Emergency Medical Service Operating Sites. In *International Workshop on Machine Learning, Optimization, and Big Data*. Springer, 293–305. https://doi.org/10.1007/978-3-319-51469-7_25
- [4] Chun-Ling Chuang and Rong-Ho Lin. 2007. A maximum expected covering model for an ambulance location problem. *Journal of the Chinese Institute of Industrial Engineers* 24, 6 (2007), 468–474. <https://doi.org/10.1080/10170660709509061>
- [5] Richard Church and Charles ReVelle. 1974. The maximal covering location problem. In *Papers of the regional science association*, Vol. 32. Springer-Verlag Berlin/Heidelberg, 101–118. <https://doi.org/10.1007/BF01942293>
- [6] Fabrice Dami, Christel Golay, Mathieu Pasquier, Vincent Fuchs, Pierre-Nicolas Carron, and Olivier Hugli. 2015. Prehospital triage accuracy in a criteria based dispatch centre. *BMC emergency medicine* 15 (2015), 1–9.
- [7] Mark S Daskin. 1983. A maximum expected covering location model: formulation, properties and heuristic solution. *Transportation science* 17, 1 (1983), 48–70. <https://doi.org/10.1287/trsc.17.1.48>
- [8] Karl F Doerner, Walter J Gutjahr, Richard F Hartl, Michaela Karall, and Marc Reimann. 2005. Heuristic solution of an extended double-coverage ambulance location problem for Austria. *Central European Journal of Operations Research* 13, 4 (2005), 325–340.
- [9] Mickey S Eisenberg, Bruce T Horwood, Richard O Cummins, Robin Reynolds-Haertle, and Thomas R Hearne. 1990. Cardiac arrest and resuscitation: a tale of 29 cities. *Annals of emergency medicine* 19, 2 (1990), 179–186. [https://doi.org/10.1016/S0196-0644\(05\)81805-0](https://doi.org/10.1016/S0196-0644(05)81805-0)
- [10] Erhan Erkut, Armann Ingólfsson, and Güneş Erdoğan. 2008. Ambulance location for maximum survival. *Naval Research Logistics (NRL)* 55, 1 (2008), 42–58. <https://doi.org/10.1002/nav.20267>
- [11] Takehiro Furuta and Ken-ichi Tanaka. 2013. Minisum and minimax location models for helicopter emergency medical service systems. *Journal of the Operations Research Society of Japan* 56, 3 (2013), 221–242. <https://doi.org/10.15807/jorsj.56.221>
- [12] Michel Gendreau, Gilbert Laporte, and Frédéric Semet. 1997. Solving an ambulance location model by tabu search. *Location science* 5, 2 (1997), 75–88. [https://doi.org/10.1016/S0966-8349\(97\)00015-6](https://doi.org/10.1016/S0966-8349(97)00015-6)
- [13] Michel Gendreau, Gilbert Laporte, and Frédéric Semet. 2001. A dynamic model and parallel tabu search heuristic for real-time ambulance relocation. *Parallel computing* 27, 12 (2001), 1641–1653. [https://doi.org/10.1016/S0167-8191\(01\)00103-X](https://doi.org/10.1016/S0167-8191(01)00103-X)
- [14] Michel Gendreau, Gilbert Laporte, and Frédéric Semet. 2006. The maximal expected coverage relocation problem for emergency vehicles. *Journal of the Operational Research Society* 57 (2006), 22–28. <https://doi.org/10.1057/palgrave.jors.2601991>
- [15] M Guida, R Calabria, and G Pulcini. 1989. Bayes inference for a non-homogeneous Poisson process with power intensity law (reliability). *IEEE Transactions on Reliability* 38, 5 (1989), 603–609. <https://doi.org/10.1109/24.46489>
- [16] Kathleen Hogan and Charles Revelle. 1986. Concepts and applications of backup coverage. *Management science* 32, 11 (1986), 1434–1444. <https://doi.org/10.1287/mnsc.32.11.1434>
- [17] Syed A Hossain and Ram C Dahiya. 1993. Estimating the parameters of a non-homogeneous Poisson-process model for software reliability. *IEEE Transactions on Reliability* 42, 4 (1993), 604–612. <https://doi.org/10.1109/24.273589>
- [18] Diane Laatz, Tyson Welzel, and Willem Stassen. 2019. Developing a South African helicopter emergency medical service activation screen (SAHAS): a Delphi study. *African Journal of Emergency Medicine* 9, 1 (2019), 1–7. <https://doi.org/10.1016/j.afjem.2018.09.001>
- [19] Karl D Majeske. 2007. A non-homogeneous Poisson process predictive model for automobile warranty claims. *Reliability Engineering & System Safety* 92, 2 (2007), 243–251. <https://doi.org/10.1016/j.res.2005.12.004>

- [20] Matthew S Maxwell, Mateo Restrepo, Shane G Henderson, and Huseyin Topaloglu. 2010. Approximate dynamic programming for ambulance redeployment. *INFORMS Journal on Computing* 22, 2 (2010), 266–281. <https://doi.org/10.1287/ijoc.1090.0345>
- [21] Joe Naoum-Sawaya and Samir Elhedhli. 2013. A stochastic optimization model for real-time ambulance redeployment. *Computers & Operations Research* 40, 8 (2013), 1972–1978. <https://doi.org/10.1016/j.cor.2013.02.006>
- [22] Amir Ali Nasrollahzadeh, Amin Khademi, and Maria E Mayorga. 2018. Real-time ambulance dispatching and relocation. *Manufacturing & Service Operations Management* 20, 3 (2018), 467–480. <https://doi.org/10.1287/msom.2017.0649>
- [23] Kobi Peleg and Joseph S Pliskin. 2004. A geographic information system simulation model of EMS: reducing ambulance response time. *The American journal of emergency medicine* 22, 3 (2004), 164–170. <https://doi.org/10.1016/j.ajem.2004.02.003>
- [24] Charles ReVelle and Kathleen Hogan. 1989. The maximum availability location problem. *Transportation science* 23, 3 (1989), 192–200. <https://doi.org/10.1287/trsc.23.3.192>, arXiv:<https://doi.org/10.1287/trsc.23.3.192>
- [25] Emmanuel S Savas. 1969. Simulation and cost-effectiveness analysis of New York's emergency ambulance service. *Management science* 15, 12 (1969), B–608. <https://doi.org/10.1287/mnsc.15.12.B608>
- [26] Haipeng Shen and Jianhua Z. Huang. 2008. Forecasting time series of inhomogeneous Poisson processes with application to call center workforce management. *The Annals of Applied Statistics* 2, 2 (2008), 601 – 623. <https://doi.org/10.1214/08-AOAS164>
- [27] Masoud Swalehe and Semra Gunay Aktas. 2016. Dynamic ambulance deployment to reduce ambulance response times using geographic information systems: A case study of Odunpazari District of Eskisehir Province, Turkey. *Procedia Environmental Sciences* 36 (2016), 199–206. <https://doi.org/10.1016/j.proenv.2016.09.033>
- [28] Constantine Toregas, Ralph Swain, Charles ReVelle, and Lawrence Bergman. 1971. The location of emergency service facilities. *Operations research* 19, 6 (1971), 1363–1373. <https://doi.org/10.1287/opre.19.6.1363>
- [29] Terence D Valenzuela, Denise J Roe, Graham Nichol, Lani L Clark, Daniel W Spaite, and Richard G Hardman. 2000. Outcomes of rapid defibrillation by security officers after cardiac arrest in casinos. *New England Journal of Medicine* 343, 17 (2000), 1206–1209. <https://doi.org/10.1056/NEJM200010263431701>
- [30] Martin Van Buuren, Rob van der Mei, Karen Aardal, and Henk Post. 2012. Evaluating dynamic dispatch strategies for emergency medical services: TIFAR simulation tool. In *Proceedings of the 2012 Winter Simulation Conference (WSC)*. IEEE, 1–12. <https://doi.org/10.1109/WSC.2012.6465214>
- [31] Jonathan Weinberg, Lawrence D Brown, and Jonathan R Stroud. 2007. Bayesian forecasting of an inhomogeneous Poisson process with applications to call center data. *J. Amer. Statist. Assoc.* 102, 480 (2007), 1185–1198. <https://doi.org/10.1198/01621450600001455>