

# Using Big Data to optimize dynamic ambulance availability maps: bridging the gap in emergency services

## Abstract

In this study, we explore the use of Big Data to dynamically optimize ambulance availability and response times in emergency medical services (EMS). Integrating Big Data principles with the Open Source Routing Machine (OSRM) routing engine, a novel algorithm was used to calculate travel times across an irregular grid, creating real-time, colour-coded time-accessibility maps. In contrast with traditional static models, this approach updates dynamically, accounting for road conditions, accidents and other disruptions to minimize delays. By excluding unusable segments from calculations, the algorithm ensures rapid recalculations, maintaining EMS coverage in evolving conditions. Testing showed significant improvements in response time estimation and resource allocation, particularly in urban environments with complex road networks. This real-time mapping tool offers EMS dispatchers an enhanced decision-support system, potentially saving lives by reducing ambulance response times and improving service availability across diverse areas.

## Keywords

Real-time mapping • Big Data • emergency medical services (EMS)

Michał Lupa<sup>1</sup>, Weronika Paterek<sup>2</sup>,  
Mateusz Zawadzki<sup>3</sup>, Michał Chromiak<sup>4</sup>,  
Katarzyna Adamek<sup>2</sup>

<sup>1</sup>Faculty of Space Technologies, AGH University of Krakow,  
Cracow, Poland  
e-mail: e-mail: mlupa@agh.edu.pl

<sup>2</sup>Department of Geoinformatics and Applied Computer  
Science, AGH University of Krakow, Cracow, Poland  
e-mail: wnaziemiec@agh.edu.pl  
e-mail: kadamek@agh.edu.pl

<sup>3</sup>Faculty of Earth Sciences and Spatial Management,  
Maria Curie-Skłodowska University in Lublin, Lublin, Poland  
e-mail: mateusz.zawadzki@mail.umcs.pl

<sup>4</sup>Department of Cybersecurity and Computational Linguistics,  
Maria Curie-Skłodowska University in Lublin, Lublin, Poland  
e-mail: michal.chromiak@mail.umcs.pl

Received: 18 December 2024

Accepted: 20 May 2025

## Introduction

The primary task of emergency services is to reach an incident scene in the shortest possible time. One of the indicators of emergency service effectiveness is the Ambulance Response Time (ART). This is defined as the time from the moment an ambulance is called until EMS arrival on scene (Castrén et al. 2008; Brice et al., 2022). The time of providing medical assistance to a patient is a key factor determining their chances of survival or maintaining health (Alrawashdeh et al. 2021; Blanchard et al. 2012; Afzali et al. 2021). Approximately 10% of fatalities from incidents occur within 3–5 min, and 60% occur within 30 min (Terzi et al. 2013). Therefore, identifying optimal EMS station locations has been a key research focus with the aim of maximizing coverage and minimizing response times. The Maximal Covering Location Model has been instrumental in efforts to station EMS units strategically, covering as large an area as possible within a given response time (Azizan et al. 2012; Li et al. 2018; Zhou et al. 2021; Li et al. 2015). Geographic Information Systems (GIS) have further supported ART improvements by helping model and visualize EMS coverage zones. In this context, the service area is the minimum time in which an ambulance can reach the incident scene or the maximum legally regulated time in which EMS must reach the patient (Piórkowski 2018). Branas et al. (2000) laid the groundwork for systematic GIS use in EMS deployment, leading to more studies examining spatial distribution and response zones. Peleg and Pliskin (2004) assessed EMS coverage in Haifa, Israel, based on 8- and 15-min ART zones. Meanwhile, research on the Odunpazari District in Turkey developed coverage maps for 5-min response times (Swalehe & Aktas 2016). In another approach, Shuib and Zaharudin (2010) used historical data to create EMS service zones, showing that past incident trends can guide future dispatch decisions. Analyzing historical data on ambulance

calls can support future dispatcher decision-making. This has been further highlighted Budge et al. (2010) and Ingólfsson (2013) who determined service area coverage based on the probability that the nearest available ambulance could respond to a call. A similar approach to availability maps was presented by Westgate (2016). Here, the probability of the nearest ambulance in Toronto arriving at the scene within 4 min was analyzed. Vanderschuren & McKune (2015) investigated ambulance availability in rural areas with a focus on the South Africa's Western Cape province. EMS service areas need to adapt dynamically to changing conditions including road repairs, traffic jams and mass events. Dynamically deployed ambulances can reduce response times effectively. Lam et al. (2015) showed that real-time ambulance relocation throughout the day improved response times. Similarly, Hashtarkhani et al. (2023) used CPLEX optimization to increase EMS coverage within a 5-minute response area, from 69% to 75%. Andersson and Vårbrand (2007) and Schmid (2012) introduced dynamic relocation models, reducing travel times by nearly 13% using Approximate Dynamic Programming. Wajid et al. (2020) used the Double Standard Model with Google Maps to more accurately match ambulance locations with accident sites. The optimal deployment of EMS units is critical for ART. Data-driven approaches that analyse historical and real-time data are particularly effective. Ji et al. (2020) demonstrated that ambulances redeployed in real time to high-demand areas further reduced ART. Lupa et al. (2021) found that analyzing travel times using actual EMS speeds provides dispatchers with a reliable decision-support tool, though complex urban routing can sometimes delay real-time calculations.

In this study, we aimed to address the limitations of traditional, static EMS coverage maps by developing a dynamic tool that adjusts in real-time to factors such as traffic incidents and road closures. By leveraging Big Data techniques and the Open

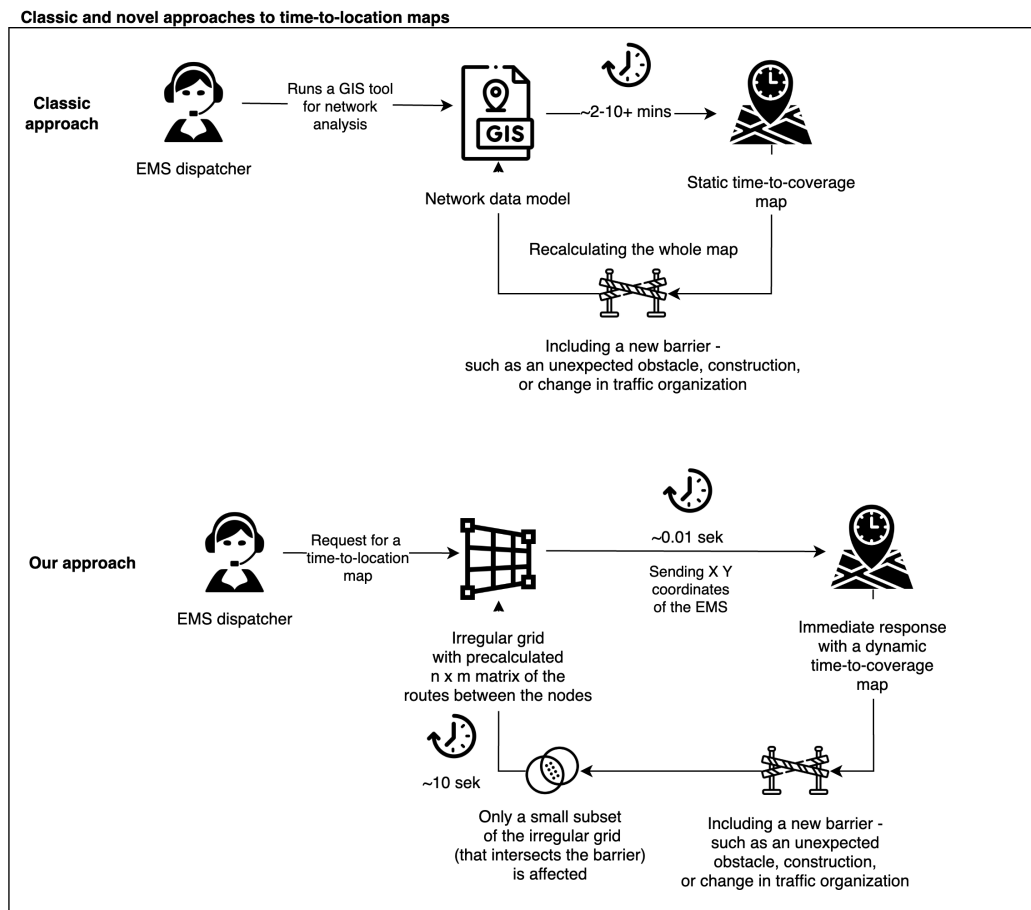


Figure 1. A schematic comparison of traditional methods for generating static time-to-location maps. This shows isochrones of travel time within a given area, and the approach proposed by the authors, which enables dynamic real-time generation of such maps  
Source: own elaboration

Source Routing Machine (OSRM) engine, this tool calculates response zones dynamically across an irregular computational grid, considering real-time road conditions and incident data. This approach offers EMS dispatchers an enhanced decision-support system, helping them optimize ART based on live conditions and potentially saving lives through faster, more efficient EMS deployment.

**Methods and dataset**  
**Design principles and system architecture**

Based on research interviews conducted with medical dispatch centre employees, the current process of dispatching an ambulance to the incident site, from the moment the call is received, is as follows:

The call is received by the intake dispatcher → A Code 1 or 2 is assigned, where 1 indicates a priority status → The system suggests the nearest ambulance based on estimated travel time. However, this estimate does not account for traffic, obstacles such as roadblocks or construction, or actual emergency speeds → The dispatcher chooses the ambulance, using system recommendations, situational awareness and personal experience → The type of call and ambulance availability in the area are considered. → If no ambulances are available, even priority calls are put on hold.

Current systems lack tools for real-time EMS coverage analysis, leading to gaps during major incidents. Ambulances

transporting patients outside their zones must return to their base before responding to new calls, which limits availability. Key bottlenecks include reliance on inaccurate travel times, which do not consider rush hours, road capacities and emergency speeds. This can potentially delay care. Even slight delays can impact survival rates and reassigning ambulances requires tools for analyzing coverage impacts. Dynamic time-to-location maps displaying ambulance coverage for an entire region would minimize gaps, enabling more effective EMS access planning.

Dynamic time-to-location maps enhance usability for critical decision-making. Minimizing map generation time is essential. Figure 1 compares classic (static) GIS-generated maps with the proposed dynamic approach. Static maps must be regenerated using GIS tools whenever an ambulance's position changes, limiting their use to base locations. They cannot dynamically define service areas ambulances can reach within a specific time (isochrone) and require full regeneration if barriers such as road closures arise. With the proposed solution, time-to-location maps are generated in real-time with minimal overhead (<1 s). Dynamic isochrone determination is based on an irregular grid and a pre-calculated travel time matrix between nodes. When road closures occur, calculations focus only on affected nodes rather than the entire matrix. This approach produces isochrones that dynamically adjust to ambulance location changes, ensuring efficient and accurate service area representation.

To enable these, the solution leverages Big Data tools, including Apache Spark, GeoTools and distributed spatial data (SRDD – Spatial Resilient DistributedDataset) stored in the Hadoop Distributed File System. Using supercomputers such as Zeus and Prometheus, the PLGrid infrastructure facilitates parallel and distributed processing across multiple machines, enabling rapid, large-scale calculations. The system includes five Ubuntu instances with 8 GB RAM, four single-core processors and essential software (Apache Spark, Hadoop, OSRM) in a cluster with four worker nodes and one master. Using Secure Shell Protocol, the Spark cluster manager co-ordinates task distribution, with a driver process on the main node assigning tasks to executors on each worker node.

**Algorithm for travel time calculation**

The algorithm builds an accurate road network by snapping a 2D uniform grid at 250-m intervals to real road nodes from OpenStreetMap, with added ambulance station points (137 in Malopolska) to create an irregular grid. Using OSRM's contraction hierarchy algorithm, routes and travel times between grid points are calculated. This forms a travel time matrix stored in a distributed SRDD file. Ambulance-specific speed profiles adjusted from GPS data (2013–2023) using a Monte Carlo approach produce a custom velocity vector (147, 122, 80, 69, 38 km/h), enabling realistic routing in emergencies.

Availability maps show coverage through the following process: requests with selected co-ordinates are sent to the controller, where the SRDD filters routes originating from those points. These routes are mapped to endpoint pairs and corresponding times. For each node, the shortest travel time is calculated, and the data is then transmitted to the view layer. The coverage map is colour-coded to indicate travel times, namely, pink for 0 min (ambulance station), green for 0–7 min, yellow for 7–14 min, red for 14–20 min and black for over 20 min.

In dynamic scenarios with temporarily unusable roads such closures following accidents or repairs, the algorithm recalculates routes excluding blocked segments to avoid delays. Instead of recomputing all the travel times (n<sup>2</sup> HTTP requests), it only updates the affected routes, which reduces the recalculation time. Using GeoSpark, the algorithm computes the fastest estimated time of arrival (ETA – Estimated Time of Arrival) from a point of interest across all nodes. It then generates colour-coded isochrones that dynamically map accessibility for changing networks.

**Performance considerations**

In static cases, a single route calculation suffices for ambulance availability maps. However, static maps become outdated with infrastructure changes, risking delays that affect patient outcomes. To mitigate these delays, in this study, we have explored architectural optimizations to accelerate map updates. Using parallelization, we assessed how temporal coverage map updates can be expedited by analyzing hardware configurations and software characteristics. Sensitivity testing focused on several key factors, that is, the size of the route dataset in the distributed system, the number of modified roads due to infrastructure changes, the number of deployed vehicles and the route determination algorithm used for calculating the fastest routes.

**Dataset**

In this study, we used data from the Emergency Notification Centre, provided as csv files organized into two schemas. The first schema detailed EMS deployments, including dispatch number, dispatch/request/return times, call reason, departure location and basic dispatch report details. A key feature was the GPS ID number, linking to detailed GPS logs. GPS data included EMS routes, vehicle speeds and timestamps, covering from

*Table 1. Average search time of the distributed dataset depending on the total number of routes stored*

Route Count (mln)	Time (s)
0.5	0.7
1.0	1.1
1.5	1.4
2.0	1.8
2.5	2.2

Source: own elaboration

2020 to 2023. The dataset was used to create an irregular grid of approximately 5,000 points, resulting in 25 million possible routes (5000<sup>2</sup> = 25,000,000). Due to storage and RAM constraints, testing focused on subsets with 0.5 million routes to evaluate the number of worker nodes in the computer cluster and 2.5 million routes to assess the impact of cores and available RAM. These subsets correspond to routes from 100 (5000 choose 100 = 0.5 million) and 500 (5000 choose 500 = 2.5 million) randomly selected points, respectively.

**The tested operations**

The tested operations for time measurement included route search finding routes intersecting user-specified roads, OSRM server restart updating exclusions for restricted roads, route retrieval (5,000 HTTP requests and storage in the distributed dataset), and temporal coverage map loading which calculated the shortest times from 150,000 routes across 30 vehicle locations.

**Results**

**Number of routes in the dataset**

The primary non-computational factor impacting programme performance is the number of routes in the distributed dataset. Conducted on four quad-core machines (each with 4 GB RAM and a single-core processor), the study encompassed a dataset with 0.5 million routes.

**Searching the route dataset**

Search time increases linearly with the number of routes and each additional 0.5 million routes adds approximately 0.4 s (Tab. 1).

**Retrieving routes**

The main time overhead in retrieving routes stems from HTTP requests, while data saving has minimal impact. Averaged measurements show that dataset volume does not significantly affect retrieval time.

**Loading the temporal coverage map**

Calculating shortest travel times adds only 0.02 s per additional 0.5 million routes, with an imperceptible impact on performance.

**Number of modified roads**

The application memory can store many modified roads that have previously been excluded. Following the exclusion of roads, the OSRM server restart must be conducted with consideration of all previously modified roads. Therefore, a study was conducted to determine whether the number of modified roads affects the time required to restart the routing engine.

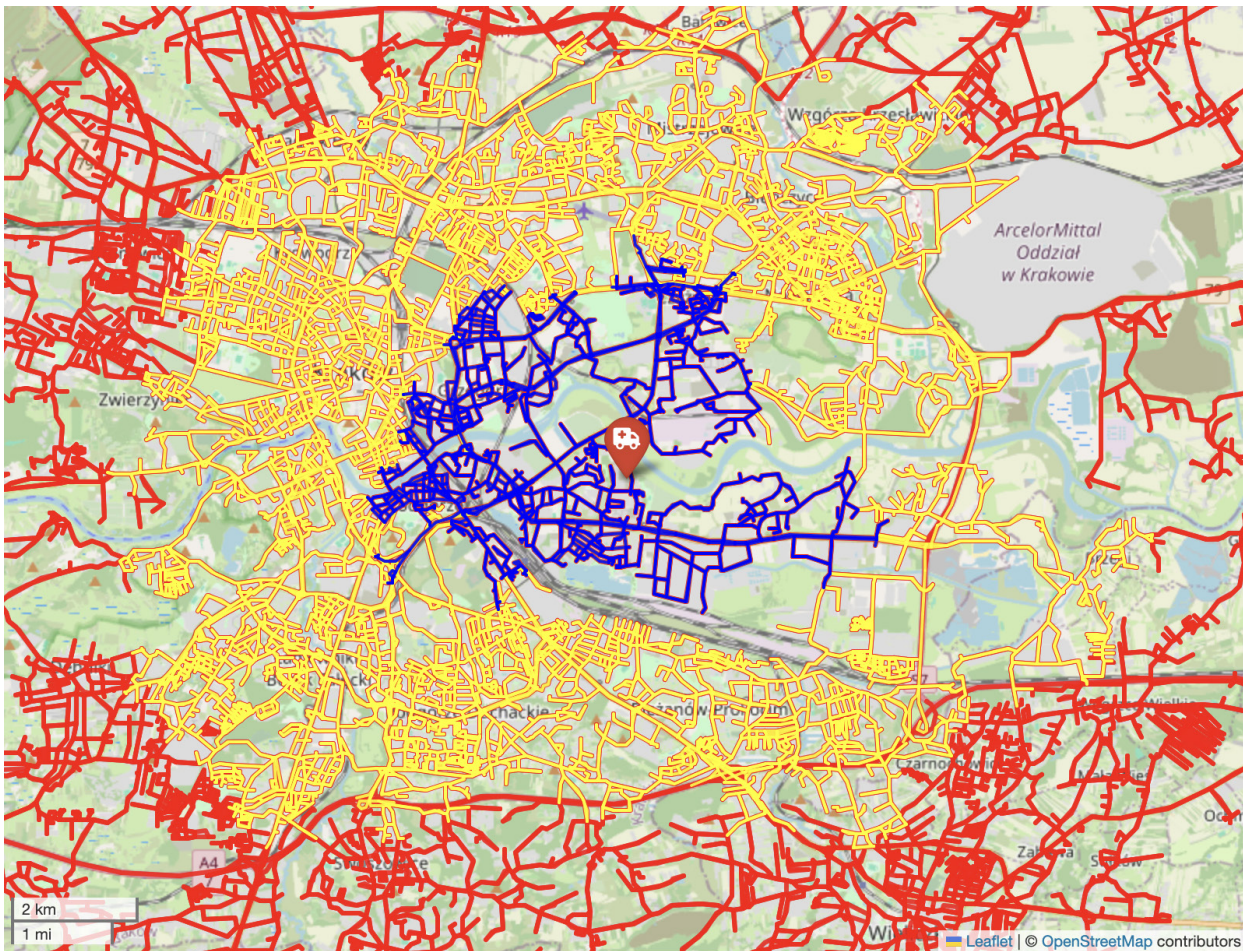


Figure 2. Temporal coverage map for an ambulance in the centre of the blue area. Blue 5 min, green: 10 min, red 15 min  
Source: own elaboration

**Restarting the routing engine**

**Number of vehicles**

The number of vehicles set by the user may only affect the speed of calculating and loading the temporal coverage map. Therefore, only this one functionality was subjected to testing.

**Loading the temporal coverage map**

The map below (Figure 2) shows an example of a dynamic time-to-location map for a single ambulance, where the grid of points used for calculations has been snapped to the road network. The colours of each segment represent the time ranges the ambulance can reach within 5, 10 and 15 min.

**Fastest route determination algorithm**

The OSRM engine provides two methods for calculating the fastest routes, namely, Contraction Hierarchies and Multi-Level Dijkstra. Each of these impact restart speed and data retrieval differently due to specific preprocessing requirements. Tests on machines with 8 GB RAM and single-core processors showed that the Dijkstra algorithm significantly reduces restart time, performing more than twice as fast as contraction hierarchies. According to OSRM developers, Dijkstra is ideal for scenarios with frequent data updates, though it may slow down on long routes (Running OSRM 2024).

**Restarting the routing engine**

After applying the Dijkstra algorithm, restarting the engine proves to be more than twice as fast with a shortener restart time compared to contraction hierarchies. This is especially valuable in environments with frequent updates.

**Retrieving routes**

Testing also assessed how switching from contraction hierarchies to Dijkstra affects retrieval speed. The Dijkstra algorithm proved faster, as shown in Figure 3, likely due to the shorter (under 30 km) routes in Krakow. This makes Dijkstra optimal for generating dynamic temporal coverage maps.

**Discussion and conclusions**

When updating the distributed dataset after excluding certain road network segments, spatial data processing methods were used to minimize the volume of newly retrieved information. The method for identifying routes requiring modification is detailed in a previous paper (Anonymous 2023). By analyzing the execution time of various operations, the benefits of applying geometric route processing can be assessed. The dataset used in the conducted studies comprised 2.5 million elements. The average time to download and save 5,000 routes was just under 11 s. Consequently, re-downloading the entire dataset of

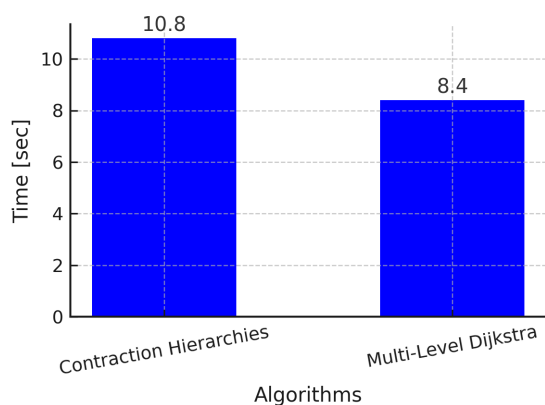


Figure 3. Average time for downloading and saving routes depending on the route determination algorithm used  
Source: own elaboration

2.5 million elements would take approximately 5,500 s (or 1.5 h). When comparing this result with the time required to update only a portion of the data, it is essential to consider the time for information retrieval and dataset searching. The anticipated time to locate specific routes is estimated at 2.2 s. In this dataset, the modified roads intersect with an average of 63,000 other roads. Therefore, the entire map updating process should take approximately 139 s (slightly over 2 min).

This outcome signifies an almost 40-fold acceleration in application performance. The most significant factor influencing the efficiency of the programme's operation is not hardware parameters but rather the structure of the algorithm and its implementation method. An algorithm designed to re-download the entire dataset from scratch does not need to retrieve and store the co-ordinates of routes between points. Retrieving

information only about travel time, without considering its route, would be somewhat faster. Comparing these two algorithmic approaches would require further research. However, employing spatial dataset searching will likely still be faster than sending queries for the Cartesian product of all grid points.

Developing applications for generating dynamic temporal coverage maps for emergency vehicles primarily focuses on the challenge of maximizing the acceleration of programme operations. This objective is achievable by implementing an irregular grid of points, selecting an appropriate route determination algorithm and applying parallel and distributed processing techniques.

Using tools such as the GeoSpark programming platform, the Open Source Routing Machine engine and PLGrid infrastructure has facilitated the creation of an application that provides dynamic temporal coverage maps and efficient parallelization. The most significant acceleration was achieved by operating on the geometries of routes previously downloaded from an external source, namely, the routing engine. Applying spatial dataset searching substantially accelerates the process of updating the temporal coverage map. An increase in the number of cluster nodes and an enhancement of their hardware parameters, such as the number of processor cores and RAM capacity, were found to impact the programme's performance. Therefore, constructing applications performing parallel computations can contribute to enhancing the efficiency of actions undertaken by various emergency services. This approach highlights the importance of leveraging advanced computational strategies and technologies to support critical operations, potentially saving lives and optimizing response times in emergency situations.

#### ORCID

Michał Lupa <https://orcid.org/0000-0002-4870-0298>

Mateusz Zawadzki <https://orcid.org/0000-0002-2126-2822>

Michał Chromiak <https://orcid.org/0000-0002-6376-808X>

Katarzyna Adamek <https://orcid.org/0000-0003-2595-9296>

#### References

- Alrawashdeh, A, Nehme, Z, Williams, B, Smith, K, Brennan, A & Dinh, DT 2021, 'Impact of Emergency Medical Service delays on time to reperfusion and mortality in STEMI', *Open Heart*, vol. 8, no. 1.
- Afzali, F, Jahani, Y, Bagheri, F & Khajouei, R 2021, 'The impact of the Emergency Medical Services (EMS) automation system on patient care process and user workflow', *BMC Medical Informatics and Decision Making*, vol. 21, no. 1.
- Andersson, T & Vårbrand, P 2007, 'Decision support tools for ambulance dispatch and relocation', *Journal of the Operational Research Society*, vol. 58, no. 2, pp. 195–201.
- Azizan, MH, Lim, CS, Hatta, WALWM & Gan, LC 2012, 'Application of OpenStreetMap data in ambulance location problem' in *2012 Fourth International Conference on Computational Intelligence, Communication Systems and Networks, Phuket, Thailand*, IEEE, pp. 321–325.
- Blanchard, IE, Doig, CJ, Hagel, BE, Anton, AR, Zygun, DA, Kortbeek, JB, Powell, DG, Williamson, TS, Fick, GH & Innes, GD 2012, 'Emergency Medical Services response time and mortality in an urban setting', *Prehospital Emergency Care*, vol. 16, no. 1, pp. 142–151.
- Branas, CC, MacKenzie, EJ & ReVelle, CS 2000, 'A trauma resource allocation model for ambulances and hospitals', *Health Services Research*, vol. 35, no. 2, pp. 489–507.
- Brice, SN, Boutilier, JJ, Gärtner, D, Harper, P, Knight, V, Lloyd, J & Tuson, M 2022, 'Emergency services utilization in Jakarta (Indonesia): a cross-sectional study of patients attending hospital emergency departments', *BMC Health Services Research*, vol. 22, no. 1.
- Budge, S, Ingolfsson, A & Zerom, D 2010, 'Empirical analysis of ambulance travel times: the case of Calgary emergency medical services', *Management Science*, vol. 56, no. 4, pp. 716–723.
- Castrén, M, Karlsten, R, Lippert, F, Christensen, EF, Bovim, E & Kvam, AM et al. 2008, 'Recommended guidelines for reporting on emergency medical dispatch when conducting research in emergency medicine: The Utstein style', *Resuscitation*, vol. 79, no. 2, pp. 193–197.
- Hashtarkhani, S, Matthews, SA, Yin, P, Mohammadi, A, Mohammad Ebrahimi, S, Tara, M & Kiani, B 2023, 'Where to place emergency ambulance vehicles: use of a capacitated maximum covering location model with real call data', *Geospatial Health*, vol. 18, no. 2.
- Ingolfsson, A 2013, 'EMS planning and management' in *International Series in Operations Research and Management Science*, ed. GS Zaric, Springer, pp. 105–128.
- Ji, S, Zheng, Y, Wang, W & Li, T 2020, 'Real-time ambulance redeployment: a data-driven approach', *IEEE Transactions on Knowledge and Data Engineering*, vol. 32, no. 11, pp. 2213–2226.
- Lam, SSW, Zhang, J, Zhang, ZC, Oh, HC, Overton, J & Ng, YY et al. 2015, 'Dynamic ambulance reallocation for the

- reduction of ambulance response times using system status management', *American Journal of Emergency Medicine*, vol. 33, no. 2, pp. 159–166.
- Li, Y, Zheng, Y, Ji, S, Wang, W, U, LH & Gong, Z 2015, 'Location selection for ambulance stations: a data-driven approach' in *Proceedings of the 23rd SIGSPATIAL International Conference on Advances in Geographic Information Systems (SIGSPATIAL '15)*, New York, NY, USA, pp. 1-4.
- Li, M, Vanberkel, P & Carter, AJE 2018, 'A review on ambulance offload delay literature', *Health Care Management Science*, vol. 22, no. 2.
- Lupa, M, Chuchro, M, Sarlej, W & Adamek, K 2021, 'Emergency ambulance speed characteristics: a case study of Lesser Poland Voivodeship, southern Poland', *Geoinformatica*, vol. 25, pp. 775–798.
- Lupa, M, Naziemiec, W, Adamek, K & Zawadzki, M 2023, 'Methodology for creating dynamic emergency vehicle availability maps', *Polish Cartographical Review*, vol. 55, pp. 24-37.
- Peleg, K & Pliskin, JS 2004, 'A geographic information system simulation model of EMS: reducing ambulance response time', *American Journal of Emergency Medicine*, vol. 22, no. 3, pp. 164–170.
- Piórkowski, A 2018, 'Construction of a dynamic arrival time coverage map for emergency medical services', *Open Geosciences*, vol. 10, no. 1, pp. 167–173.
- Running OSRM 2024, GitHub. Available from: <<https://github.com/Project-OSRM/osrm-backend>>. [20 September 2024].
- Schmid, V 2012, 'Solving the dynamic ambulance relocation and dispatching problem using approximate dynamic programming', *European Journal of Operational Research*, vol. 219, no. 3, pp. 611–621.
- Shuib, A & Zaharudin, ZA 2010, 'Framework of tazopt model for ambulance location and allocation problem', *World Academy of Science, Engineering and Technology*, vol. 70, pp. 678–683.
- Swalehe, M & Aktas, SG 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*, vol. 36, pp. 199–206.
- Terzi, O, Sisman, A, Canbaz, S, Dünder, C & Peksen, Y 2013, 'A geographic information system-based analysis of ambulance station coverage area in Samsun, Turkey', *Singapore Medical Journal*, vol. 54, no. 11, pp. 653–658.
- Vanderschuren, M & McKune, D 2015, 'Emergency care facility access in rural areas within the golden hour?: Western Cape case study', *International Journal of Health Geographics*, vol. 14, no. 1.
- Wajid, S, Nezamuddin, N & Unnikrishnan, A 2020, 'Optimizing ambulance locations for coverage enhancement of accident sites in South Delhi', *Transportation Research Procedia*, vol. 48, pp. 280–289.
- Westgate, BS, Woodard, DB, Matteson, DS & Henderson, SG 2016, 'Large-network travel time distribution estimation for ambulances', *European Journal of Operational Research*, vol. 252, no. 1, pp. 322–333.
- Zhou, Z, Xu, Z, Liu, A, Zhou, S, Mu, L & Zhang, X 2021, 'Mapping the accessibility of medical facilities of Wuhan during the COVID-19 pandemic', *ISPRS International Journal of Geo-Information*, vol. 10, no. 5, p. 318.