Towards Building a Navigation Profile for Emergency Response Vehicles by Analyzing Incident Callout Data

Moritz Aberle
City of Freiburg im Breisgau*
moritz.aberle@freiburg.de

ABSTRACT

The spatial coverage of units and resources is a key planning metric for emergency service operations. When computing isochrones for this purpose, the results depend heavily on the utilized navigation profile, which defines vehicle-specific parameters such as speed and routing restrictions. However, there does not yet exist a model for the specific driving characteristics of emergency vehicles. We analyzed historical incident callout data in the city of *Freiburg im Breisgau* to work towards developing such a model.

Keywords

Disaster Response, Deployment Optimization, Emergency Vehicle, Navigation Profile, Routing

INTRODUCTION

The coverage problem for resource allocation is one of the key performance indicators for modern fire services (Eslamzadeh et al. 2022). The biggest contributing factor to the spatial coverage of a fire station is the time en-route, which during planning can be approximated using isochrones. These depict the geographic areas accessible within a specified travel time from a given starting point. Computing isochrones requires a road network and a navigation profile, specifying the characteristics such as speed, road usage restrictions and route choice for a particular vehicle. Open-source routing software, for example GraphHopper and OpenRouteService, ship with a variety of vehicle types, e.g., bike, car, or truck. Established profiles, however, do not particularly represent the characteristics and privileges (e.g., exceeding speed limits, disregarding certain traffic signals) of emergency vehicles. As shown in Figure 1, a car-specific isochrone covers areas, where the measured time en-route to actual incidents is larger. By analyzing historical callout data, we propose the development of a specialized navigation profile tailored to emergency vehicles. Our aim is to enhance the accuracy of coverage maps, ensuring that calculated travel times are more closely aligned with real-world observations.

RELATED WORK

Optimizing emergency vehicle deployment through travel time prediction is an established approach. Over 50 years ago, researchers studied this case with the fire department in New York City (Kolesar 1975). The model used was relatively simple, taking into account the covered area, the number of units, as well as the rate and duration of alarms. The utilized driving characteristic for fire trucks was based on earlier research using a manual process with logbooks and stopwatches (Kolesar and Walker 1974).

Woelki et al. explored routing alternatives for emergency vehicles using Dijkstra's shortest-path algorithm (Woelki et al. 2015). They argue to minimize both distance and time en-route while working to find a ranking and theoretical Pareto optimum for paths to the incident. Another study from Germany compared three different navigation systems (Van Mark et al. 2024). Among those was a novel application specialized in emergency routing, which they found to output the fastest travel times. However, the study's findings are severely limited since it did not verify if the

---

[*] Acknowledging Dr. Philipp Golecki and Martin Kupper for their support with data acquisition. This work was supported by the German Federal Ministry of Research, Technology and Space (BMFTR) under grant number 13N16812.

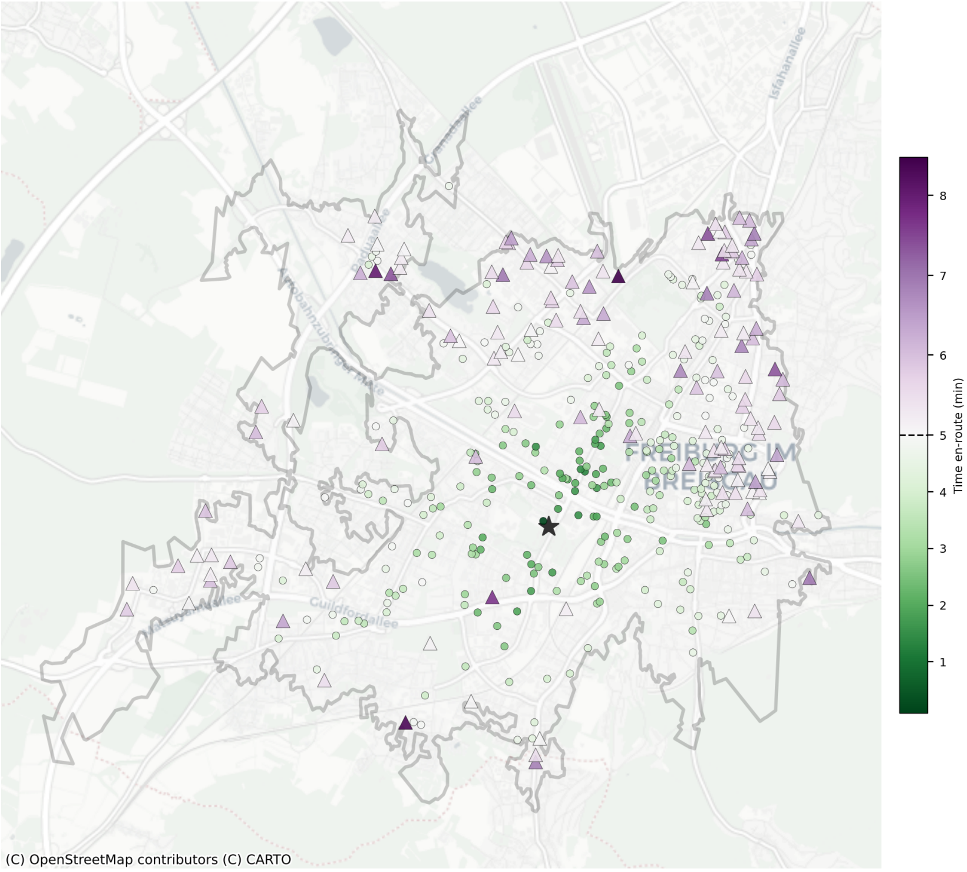

Figure 1. Incidents colored by actual time en-route within a computed five-minute isochrone, where triangles indicate a longer travel time than estimated by car-specific routing

reported durations where actually feasible in real-world settings. Potentially, the traditional systems were simply more conservative and accurate in their estimates, for example by taking into account real-time traffic conditions.

Modern studies rely on data recorded from *global navigation satellite systems* (GNSS). A long-term study conducted with GNSS recordings between 2011 and 2017 in Poland analyzed the speed of ambulances (Lupa et al. 2021). They concluded that lights-and-sirens responses are, on average, 56 % faster than normal driving. Furthermore, a correlation between the speed during those trips and the road density as well as building area could be identified. A case study from Brunswick, Germany, collected position data of 24 emergency vehicles over a five year period (Bieker-Walz et al. 2019). The analysis showed that most trips with lights-and-sirens could drive with an average speed higher than 50 km/h. Also, they identified travel time and the number of signalized intersections as factors influencing the route choice, but did not propose a navigation profile based on these findings. Similarly, Schuhmann and Lienkamp conducted a study with the Munich fire department (Schuhmann and Lienkamp 2025). The GNSS locations of five vehicles on two fire stations were recorded and a discrete route choice model was trained on this data. Their results show that one-lane streets or traveling against direction of traffic are avoided, and that residential streets or indirect routes as well as long durations make a path less attractive. On top of that, the paper describes the average speed of fire trucks, to which we will compare our results below.

The travel duration of emergency services was also evaluated in Sweden (Hassler et al. 2024). In their paper, the authors compare the actual response time to an estimation based on the fastest-path from station to scene. They find that in over 81% of the cases, the computed travel duration was underestimated, meaning the actual response time was longer. Another study from Sweden used machine learning to estimate ambulance travel times (Abid et al. 2024). Likewise Zhao and Vanberkel experimented with artificial neural networks and compared them to other approaches using data from Nova Scotia (Zhao and Vanberkel 2025).

Most notably, Zerbst et al. reported on the 2016 installation of a traffic signal priority system in our case study's region (Zerbst et al. 2016). The authors introduce a metric, the percentage of incidents where the time en-route is less than eight minutes, and examine the impact of adopting the system in specific focus areas. While effects on the response time for historical callouts are presented, no concrete values for change in vehicle speeds or conclusions regarding spatial coverage are drawn.

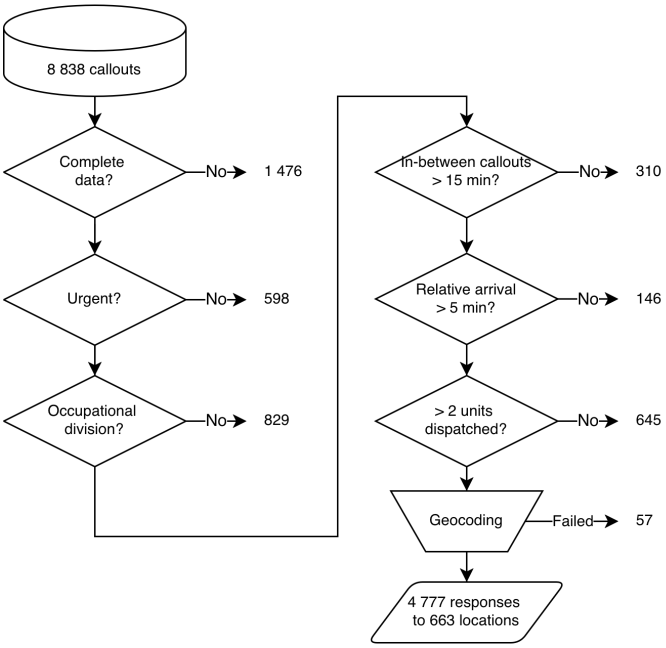

Figure 2. The pre-processing pipeline showing the number of excluded callouts in each step.

METHODOLOGY

Our analysis is based on historical callout data from the year 2024 of the fire department in the city of *Freiburg im Breisgau*, Germany. It most notably lists the date and time, street address, division, a flag for urgency, and several timestamps for every emergency response vehicle involved in each incident: T^{ack} acknowledged mission, T^{dep} departed station, T^{arr} arrived on scene, T^{bas} back in service, and T^{ret} returned to station.

During pre-processing, we excluded all non-urgent calls and focused on the occupational division (rather than volunteers) of the fire department. Furthermore, we accounted for consecutive incidents, where a vehicle would not return to its station, and removed missions where the gap between deployments was less than 15 minutes ($T_n^{\text{dep}} < T_{n-1}^{\text{ret}} + 15\text{min}$). The timings originate from the driver pushing a button on their radio unit to transmit a codified number to the dispatch center. To detect erroneous timestamps, we made the assumption that a platoon with at least two vehicles travels the same route in a convoy from station to scene, accounting for the width, height, and weight restrictions of its largest truck. Thus, any incidents where less than two units where dispatched as well as records, where the time between individual vehicle arrival on scene was greater than five minutes, were discarded.

We calculated the travel duration by subtracting the time a vehicle reported arriving on scene from when it left the station ($T^{\text{arr}} - T^{\text{dep}}$). For every location with multiple incidents (e.g., due to automatic fire alarm systems), a median arrival time was calculated to mitigate traffic congestion effects. Then, we geocoded the street address to spatial coordinates (while removing callouts where this failed) and calculated the linear distance between the vehicle's assigned station and the incident location from these. As shown in Figure 2, this finally resulted in a total of 4 777 lights-and-sirens responses of emergency vehicles to 663 unique incident locations.

Using a local GraphHopper instance (version 11.0), an open-source routing engine based on crowd-sourced OpenStreetMap data, we computed the shortest driveable route *likely* traveled by the emergency vehicle. This path consists of several segments (think from one intersection to another), including additional information such as section length, road classification, street name, as well as posted maximum speed, weight and height. An average speed for the vehicle was determined based on the distance of the (presumed) route driven and the travel time calculated above.

Finally, by aggregating all trips that pass through a particular street, an average speed was determined for this road segment. This further allowed us to group these segments by road classification.

RESULTS AND DISCUSSION

In the following section, we first describe the results obtained from analyzing vehicle callouts and the aggregation of trips on road segments. Then we propose a navigation profile and evaluate the error.

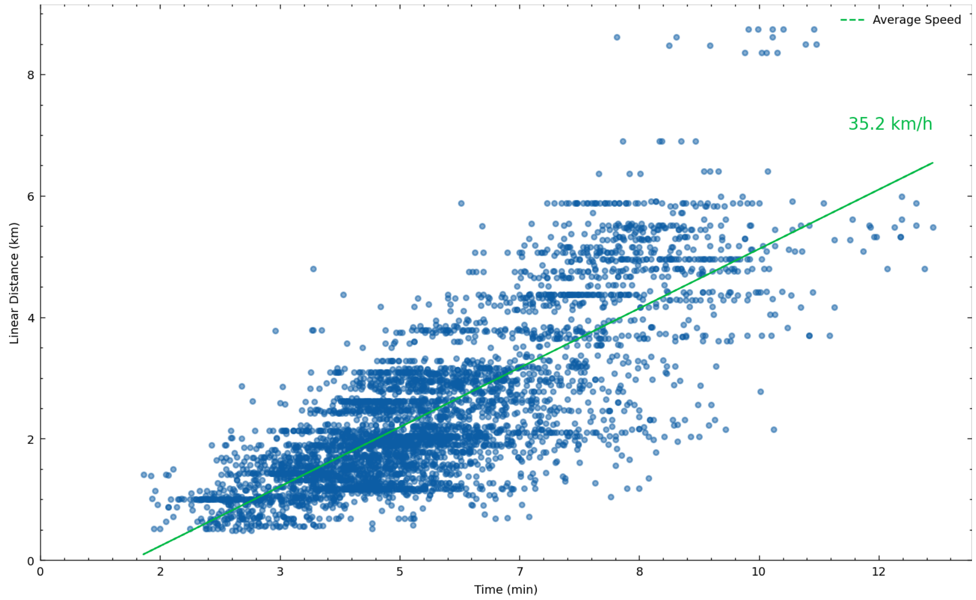

Figure 3. Scatterplot of time en-route and linear distance from station to scene

Vehicle-based Analysis

The average time elapsed driving from station to scene was 5.3 minutes, with a median of 5.0 minutes. Across all vehicles, the cumulative duration driven was 421.5 hours. Figure 3 plots the elapsed time vs. the straight-line distance from station to incident location. As a first estimate, the average speed can be fixed at 35.2 km/h.

When accounting for actual route lengths rather than distances as the crow flies, the median speed was 34.7 km/h and the mean 35.4 km/h. This is a lower bound, as it is based on the shortest possible path calculated retrospectively (median 3.1 km, mean 3.3 km) and in reality the distances traveled might be longer (resulting in larger speeds). However, it is quite consistent with GNSS data observed for fire trucks in Munich (Schuhmann and Lienkamp 2025). To compare our results in detail, we similarly split each day into night (20:00–05:59), midday (10:00–15:59), and peak (06:00–09:59 & 16:00–19:59) hours. The average calculated speed was 34.7 km/h (Munich: 30.7) at night, 35.6 km/h (30.8) midday and 33.4 km/h (30.6) during peak times. Figure 5 shows a box plot and the average speed per hour of the day. Lower average speeds are particularly observed at around 8 and 17 o'clock, likely due to rush hour and traffic congestion effects.

When comparing the time en-route observed in our data with the duration calculated by the routing engine, we find a mean absolute error (MAE) of 63.8 seconds (standard deviation: 53.5 seconds). The travel time to each unique incident location is underestimated (i.e., time en-route is actually longer than calculated) in 70.3 % of routes when using a car profile. This indicates that the resulting isochrones are likely oversized, as illustrated in Figure 4.

Road-based Evaluation

Based on the aggregation of trips on road segments, we can explore the influence of road types on speed. Using the classification from OpenStreetMap,[1] we find a median speed of 41.2 km/h on roads designated as trunk, 34.3 km/h on primary, 35.6 km/h on secondary, 37.1 km/h on tertiary, and 34.3 km/h on residential streets. For comparison with results obtained by Schuhmann and Lienkamp, we similarly group these classes into highway (motorway, trunk, primary), main (secondary, tertiary), and refer to all others as residential. Over all trips made, 23.2 % of the road segments were categorized as highways, 60.5% as main, and 16.3% as residential. The median speed calculated on highways was 36.2 km/h (Munich: 52.0 – 66.6 km/h), on main streets 35.7 km/h (43.5), and for residential roads 33.8 km/h (26.7 – 34.5). This contrasts sharply with GNSS results surveyed by others, which provides higher-resolution data, while our methodology only yields average speeds for the entire route.

A plot of all road segments colored by their average speed is shown in Figure 6. Notably, slower speeds occur mainly in the city center, while higher average speeds are found on suburban and peripheral roads. This aligns with the analysis by Lupa et al., which also identified a negative impact of building and road density on emergency vehicle speed.

---

[1] `https://wiki.openstreetmap.org/wiki/Key:highway` (last accessed February 20, 2026)

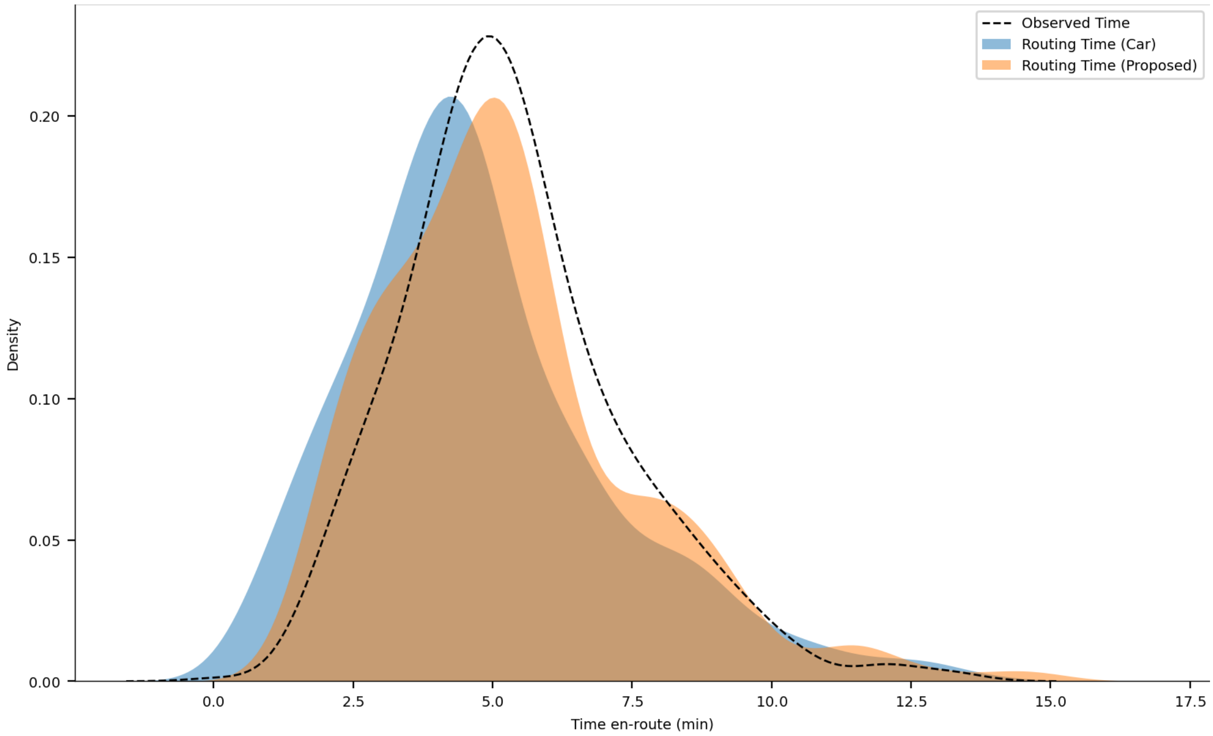

Figure 4. Kernel density estimate comparing the actual and computed times en-route to every unique incident location

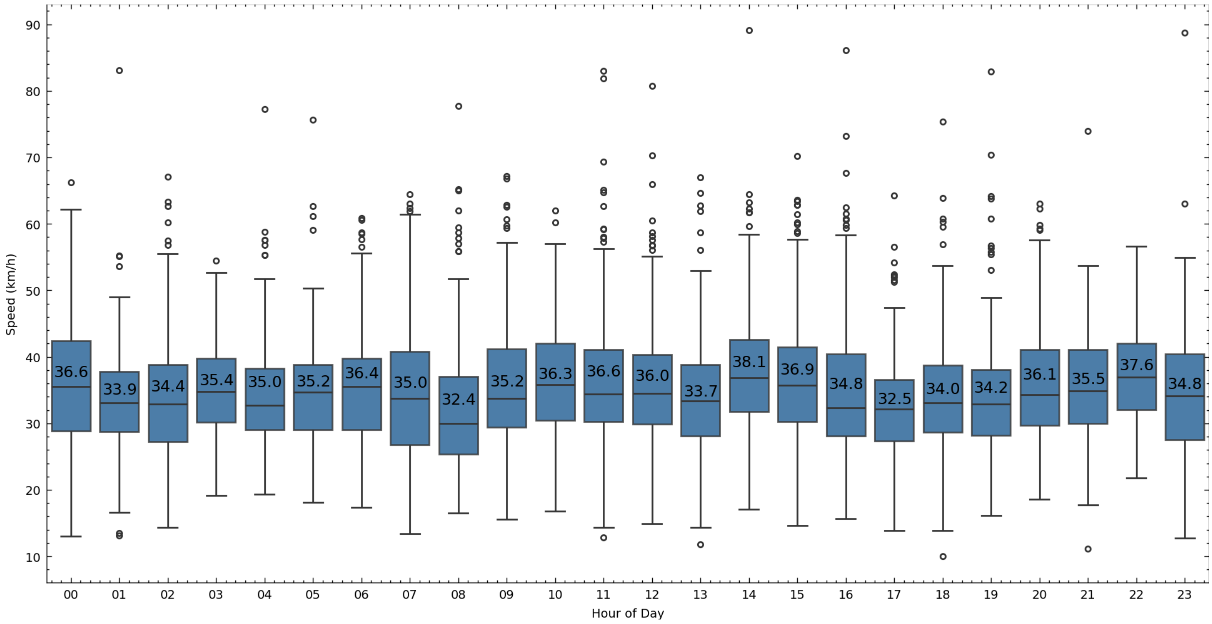

Figure 5. Box plot of the estimated driving speed for each hour of the day

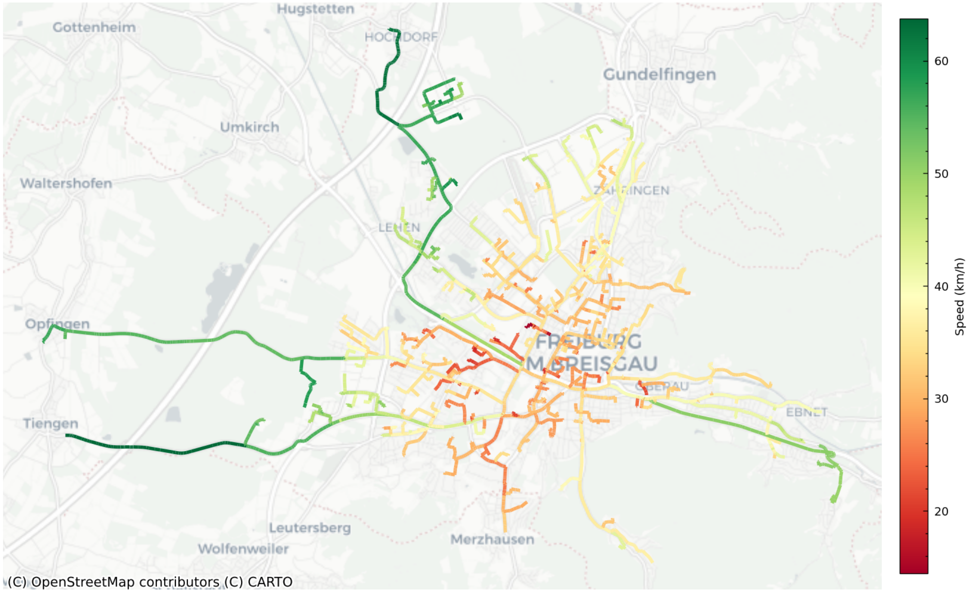

Figure 6. Road segments colored by their average speed

Navigation Profile

To demonstrate our approach, we built a navigation profile for the GraphHopper routing engine shown in Figure 8. According to the documentation,[2] the routing engine evaluates $\text{edge\_weight} = \frac{\text{edge\_distance}}{\text{speed} \cdot \text{priority}} + \text{edge\_distance} \cdot \text{distance\_influence} + \text{turn\_penalty}$ to calculate weights for all road network edges.

The *distance_influence* parameter indicates the priority given to distance over time, which is why a very small value is deliberately chosen (line 2). A *turn_penalty* of 1 000 is applied when the road class changes and the next segment is either pedestrian-only or a destination street to prioritize routing on main roads (lines 4 – 7). Restrictions regarding the vehicle dimensions and road access are added by choosing appropriate *priority* values for route choice. The maximum width, height, and weight (line 36 – 43) are specified based on values derived from the vehicles used in this case study. A malus of 0.001 is applied to all roads not accessible by cars to ensure they are avoided whenever possible (lines 44 – 47).

To find suitable *speed* parameters, we used an automatic hyperparameter optimization approach (Optuna, version 4.8.0) to minimize the mean absolute error between the routing and observed time. As a baseline configuration, the median speeds per road type from above were utilized. The resulting speeds were 74 km/h on trunk, 35 km/h on primary, 40 km/h on secondary or tertiary, 24 km/h on residential, and 32 km/h on all other road classes (lines 10 – 29). The statement of a general speed limit is inexplicable required by GraphHopper and thus set to 80 km/h (lines 30 – 33). In future work, these parameters can be detailed further, such as specifying speeds based on urban density or other edge attributes.

We again compared the actual driving duration to the time calculated by the routing engine, resulting in a MAE of 57.1 seconds (standard deviation: 56.8 seconds). The travel time is now only underestimated in 54 % of routes when using the specific emergency vehicle settings, representing an improvement of 15.3 percentage points compared to the original profile. As depicted in Figure 4, the time en-route calculated with the proposed profile more closely matches the real world timings.

Finally, we evaluate the resulting five-minute isochrone shown in Figure 7 with the car-specific coverage area in Figure 1 from the introduction. Originally, 430 incident locations were inside the isochrone, while 33.3 % of those could actually not be reached in that time. With the emergency vehicle profile, 358 unique locations fall within the coverage area, of which only 22.0 % are misclassified (i.e., they lie inside the isochrone, but the observed time en-route is actually longer).

LIMITATIONS

We did not obtain GNSS positions from vehicles and thus retrospectively estimated the path and exact distance traveled using shortest-path routing on the road network. Additionally, the data recorded in our case study describes

---

[2] `https://github.com/graphhopper/graphhopper/blob/11.x/docs/core/custom-models.md` (last accessed April 15, 2026)

an urban region. This limits the interpretability of our results: The road-based evaluation thus may not accurately portray the route choice of emergency vehicles. However, our findings are plausible and hold up to comparison with other studies using GNSS tracks.

Our data may also be subject to human bias, since the reported timings are based on manually submitted codified status reports via radio. We address this issue by only considering incidents where two or more cars from a platoon were dispatched and removing outliers with unrealistic driving durations or regarding the time between individual vehicle arrival on scene.

CONCLUSION

Compared to other approaches, our method is fast and immediately applicable, as it does not rely on GNSS data that must first be surveyed. By only using timestamp information likely already recorded by most emergency services, it can easily be applied to other regions. However, its insights are more limited, since it only captures the average speed per route, which itself is reconstructed between the station and scene.

We are certain that refining our basic navigation profile for emergency vehicles will further improve results and the applicability to other regions. As future work, the profile could be improved by GNSS data recorded from fire trucks in our case study's region over a period of six months. This enables the exact reconstruction of routes actually traveled at the respective speeds. By taking into account historical traffic data, we could also incorporate road closures and congestions. The final goal is to develop a robust routing profile that enables a realistic representation of travel times for emergency vehicles, which we would like to contribute to open-source routing engines to aid with performance optimization of fire services and during disaster recovery.

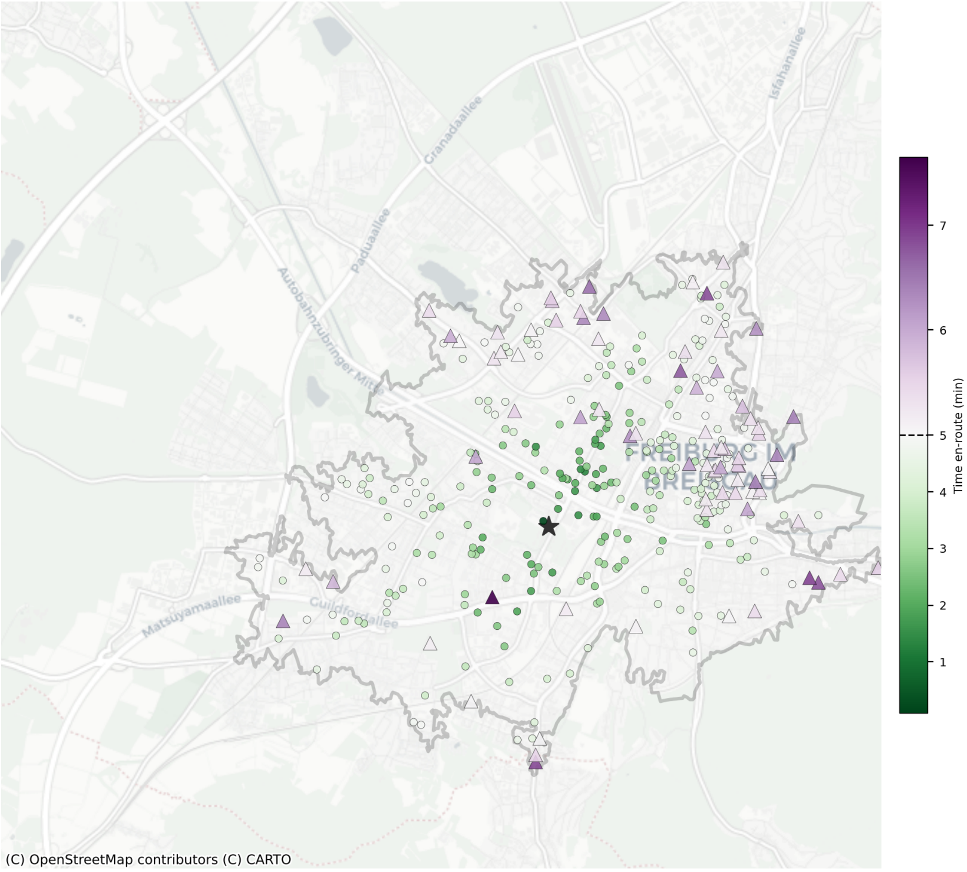

Figure 7. Incidents colored by actual time en-route within a computed five-minute isochrone using the emergency vehicle specific routing profile

REFERENCES

Abid, M. A., Lorig, F., Holmgren, J., and Petersson, J. (2024). “Ambulance Travel Time Estimation using Spatiotemporal Data”. In: *Procedia Computer Science* 238, pp. 265–272.

Bieker-Walz, L., Flotterod, Y.-P., and Sohr, A. (June 2019). “Modelling of Emergency Vehicles Route Choice with Use of Trajectory Data”. In: *2019 6th International Conference on Models and Technologies for Intelligent Transportation Systems (MT-ITS)*. Cracow, Poland: IEEE, pp. 1–8.

Eslamzadeh, M. K., Grilo, A., and Espadinha-Cruz, P. (July 27, 2022). “A Framework for Resource Allocation in Fire Departments: A Structured Literature Review”. In: *Fire* 5.4, p. 109.

Hassler, J., Andersson Granberg, T., and Ceccato, V. (Jan. 2024). “Socio-spatial Inequities of Fire and Rescue Services in Sweden: An Analysis of Real and Estimated Response Times”. In: *Fire Technology* 60.1, pp. 193–212.

Kolesar, P. (1975). “A Model for Predicting Average Fire Engine Travel Times”. In: *Operations Research* 23.4, pp. 603–613.

Kolesar, P. and Walker, W. (Jan. 1, 1974). *Measuring the Travel Characteristics of New York City’s Fire Companies*.

Lupa, M., Chuchro, M., Sarlej, W., and Adamek, K. (Oct. 1, 2021). “Emergency ambulance speed characteristics: a case study of Lesser Poland voivodeship, southern Poland”. In: *GeoInformatica* 25.4, pp. 775–798.

Schuhmann, F. and Lienkamp, M. (Sept. 8, 2025). “Understanding the Route Choice Behavior of Emergency Vehicles: A Case-Study in Munich”. In: *2025 9th International Conference on Models and Technologies for Intelligent Transportation Systems (MT-ITS)*. Luxembourg: IEEE, pp. 1–6.

Van Mark, A., Hallstein, T., Holzgreve, F., Groneberg, D. A., and Ohlendorf, D. (July 2024). “How do different navigation systems affect emergency response time? A prospective simulation study”. In: *BMJ Open* 14.7, e079094.

Woelki, M., Lu, T., and Ruppe, S. (June 2, 2015). “Ranking of alternatives for emergency routing on urban road networks”. In: València, Spain, pp. 591–598.

Zerbst, L., Vetter, M., and Koch, C.-F. (Nov. 1, 2016). “Einführung eines Systems zur Bevorrechtigung von Einsatzfahrzeugen der Feuerwehr zur Verbesserung der Hilfsfristen im Stadtgebiet Freiburg”. In: *KN - Journal of Cartography and Geographic Information* 66.6, pp. 318–324.

Zhao, Q. and Vanberkel, P. (May 16, 2025). “A Comparison of Ambulance Travel Time Approximation: Using Google Maps and Machine Learning”. In: *Proceedings of the International ISCRAM Conference*.

```json
{
  "distance_influence": 1,
  "turn_penalty": [
    {
      "if": "prev_road_class != road_class && (road_class == PEDESTRIAN ||
          ↪ road_access == DESTINATION)",
      "add": "1000"
    }
  ],
  "speed": [
    {
      "if": "road_class == TRUNK",
      "limit_to": "74"
    },
    {
      "else_if": "road_class == PRIMARY",
      "limit_to": "35"
    },
    {
      "else_if": "road_class == SECONDARY || road_class == TERTIARY",
      "limit_to": "40"
    },
    {
      "else_if": "road_class == RESIDENTIAL",
      "limit_to": "24"
    },
    {
      "else": "",
      "limit_to": "32"
    },
    {
      "if": "true",
      "limit_to": "80"
    }
  ],
  "priority": [
    {
      "if": "max_width < 3 || max_height < 3.5",
      "multiply_by": 0
    },
    {
      "if": "max_weight < 12 && max_weight_except == MISSING",
      "multiply_by": 0
    },
    {
      "if": "!car_access",
      "multiply_by": 0.0001
    }
  ]
}
```

Figure 8. An initial concept of the navigation profile *truck_emergency.json* for the GraphHopper routing engine.