

Towards a Unified Understanding of Data-Driven Support for Emergency Medical Service Logistics

Melanie Reuter-Oppermann
Karlsruhe Institute of Technology
melanie.reuter@kit.edu

Clemens Wolff
Karlsruhe Institute of Technology
clemens.wolff@kit.edu

Abstract

Time-critical medical emergencies challenge emergency medical service (EMS) systems worldwide every day. In order to respond to these incidents as soon as possible, EMS logistics' approaches can help locating and dispatching ambulances. Many of these approaches use estimates for the demand as well as the driving, service and turnaround times. In order to determine useful solutions and make informed decisions, reliable forecasts are necessary that take the characteristics and constraints of the planning problems at different levels into account. While many different approaches have been presented and tested in literature, a common understanding is still missing. This paper therefore proposes a taxonomy on EMS forecasting that distinguishes between medical emergencies and patient transports, demand and time intervals in the response process, as well as the three planning levels strategic, tactical and operational. In addition, an illustrative example and a research agenda are presented based on the findings for the taxonomy.

1. Introduction

In case of a medical emergency, patients should receive medical care as soon as possible. In many countries worldwide, emergency medical service (EMS) systems exist that send first responders, ambulances staffed with paramedics, or emergency doctors, to the patient for a first treatment and potential transport to hospital. EMS providers try to prepare for medical emergencies as good as possible, locating their ambulances in such a way that response times to patients are minimised. Unfortunately, the actual demand is not known in advance. Consequently, reliable forecasts are necessary to predict the expected demand and prepare accordingly. While in the past often only few data points were collected, more and more data becomes available now. EMS services are tracking served emergencies and transports more thoroughly, leading to more and better

data. Additionally, external data sources are or will be usable, such as weather or population data.

Due to the importance of the topic, a range of publications exist that present demand forecasting approaches. Recently, Steins et al. have investigated the influence of different parameters like population, income, work places, road traffic or weather conditions for three Swedish EMS regions [1]. While they were able to show the influence of different parameters, their rather complex approach could not significantly outperform the simpler approach that is currently used by the EMS provider. Further research is necessary to deeply investigate the relevance of different input parameters and the performance and requirements of different approaches. A common understanding of data-driven support for EMS logistics is a requirement for a thorough analysis.

In this paper an overview of existing approaches for EMS demand forecasting together with a taxonomy is presented. The taxonomy is based on three dimensions, the tasks, the entities to be forecasted and the planning levels. We discuss the requirements with relation to the planning problems for which the forecasts will be used as input. To study and validate our hypotheses and findings, we analyse an exemplary use case with EMS data from a German EMS region and present a set of research questions that will be investigated in future work.

The paper is structured as follows. In Section 2 the relevant literature on EMS logistics and EMS forecasting is summarised. Section 3 presents the derived taxonomy for data-driven support with three dimensions. In Section 4 initial results for strategic demand forecasting with machine learning approaches are presented and an outline for future research is proposed. The paper finishes with a conclusion in Section 5.

2. Literature

In this section, the relevant literature on EMS logistics and forecasting approaches, mainly from the area of operations research, is presented.

2.1. EMS logistics

Similar to the German EMS system, many EMS systems worldwide are responsible for two main tasks. First, they attend time-critical medical emergencies, e.g. by sending paramedics in ambulances. Second, patients are transported by ambulances to, from and between hospitals and other care institutions. While medical attention is necessary during the transports, they are usually not time-critical and can therefore be scheduled if known in advance. Due to various reasons, not all transports are announced in advance, but must be fulfilled on short notice.

Emergency medical service systems face many logistical problems that they need to solve regularly. An overview of EMS logistics can be found in [2]. The problem of locating ambulances and bases is the one that has been studied the most. Already in 1971 Toregas et al. proposed the Location Set Covering Model [3]. The Maximum Expected Covering Location Problem (MEXCLP) introduced by Daskin acknowledges the fact that ambulances are not always available but might be already busy serving another incident [4]. The MEXCLP is still widely used in practice and it has been applied and extended in many publications, for example by including survival probabilities [5]. Although the location of bases is actually a strategic and the location of ambulances more a tactical decision, these problems are often considered simultaneously. On the operational level, the corresponding planning problem is about the relocation of ambulances. The main aim of relocating ambulances is to increase coverage and reduce response times in real-time. Gendreau et al. were one of the first to address the real-time ambulance location problem [6]. In a simple version of the problem the decision to be made is which base an ambulance to send to after having served an emergency. In other approaches, relocations are predetermined based on time stamps or scenarios. That means that at specific points of time or when certain scenarios are met, a predefined set of ambulances is sent from its current locations to other bases. In the most sophisticated relocation approaches moves are determined in real-time, for example by solving efficient heuristics [7]. An overview of ambulance location and relocation models can be found in [8].

In many EMS systems still the closest idle ambulance is dispatched, even though Carter et al.

already showed that this approach is not always optimal for the system [9]. Dean studied this dispatching rule and found that implementing other rules could improve the overall response time performance [10]. Therefore, several approaches have been proposed, for example by Andersson and Vårbrand for low-priority calls [11]. Jagtenberg et al. have presented a dispatching heuristic that significantly reduced the fraction of late arrivals in their study with the drawback of an increased average response time [12]. Belanger et al. present recent literature and trends on location, relocation and dispatching approaches for ambulances [13].

Staff planning for EMS mainly differs between staff in the coordination centre and in the ambulances. Approaches for the ambulances rostering problem have been presented by Bradbeer et al. [14] and Rajagopalan et al. [15], for example. In coordination centres, assuming a certain call rate, it must be determined how many call takers are necessary to answer those calls. Besides staffing, task handling is an important topic. The question is if it is more efficient to divide or combine the tasks of call taking and ambulance dispatching. Van Buuren et al. use discrete-event simulations to investigate this question [16].

Regarding patient transports, the announcement of tasks as well as the policies applied in practice define if transports are scheduled in advance or assigned to ambulances in real-time throughout the day, potentially immediately before the transport needs to take place. Parragh et al. introduced formulations and solution approaches for the static patient transportation problem with different types of ambulances [17]. If all patient transport tasks are known in advance, e.g., the night before, and if the problem size allows for a solution in acceptable time, a mathematical formulation can be solved to optimality. If this is not the case, heuristics are necessary in practice. If none or only part of the tasks are known in advance and the rest becomes known throughout the day, scheduling the tasks resembles an online problem. Ardekani et al. presented three heuristics for the patient transportation problem: an insertion-based heuristic that determines a starting solution, an improvement heuristic and a dynamic heuristic that inserts short-term demand into the existing schedule in real-time [18]. For the online case, but when rescheduling is possible, Kergosien et al. proposed a tabu search heuristic, which is called every time a new transport emerges [19]. Schilde et al. studied the problem of incorporating unknown but expected return transports for the patients [20]. They model it as a dynamic stochastic problem and propose four advanced variants of metaheuristics.

Many papers study either emergency rescues or

patient transports, only few investigate both problems simultaneously. Kergosien et al. built a generic discrete event simulation-based analysis model that studies the management of a fleet of ambulances with the aim of optimally serving emergency requests as well as transporting patients between their homes and hospitals [21]. A simulation of the EMS system can be helpful to get realistic estimates on how the decisions would influence the real system. There are two main ways for the use of simulation in the decision process. Often, it is used as a stand-alone evaluation tool that compares different scenarios. Alternatively, simulation can also be used as a subroutine in the optimisation approach. A review on simulation studies for EMS systems was published by Aboueljine et al. [22].

2.2. EMS forecasting

This section gives a general overview of forecasting approaches presented in the literature, mainly for emergency rescues. A more thorough summary can be found in [2] or [23], for example.

Emergency demand forecasting

Several methods divide the year into seven days with 24 hours each or a week in 168 one hour intervals and accumulate the historical emergency demand during each hour of the week for a set of demand points, e.g. postal code regions, as for example proposed by Matteson et al. [24] or Steins et al. [25]. The forecasted number for each hour is then used as input for determining the necessary number and location of ambulances in each hour. Seasonal patterns of demand for EMS has been shown by Vile et al., for example [26] and an underlying increase over the past by Lowthian et al. [27]. To cope with fluctuations different approaches including linear, sinusoidal and support vector regression [28], simple moving averages and more complex time series approaches that include the previous hours in the forecast [29, 24] have been proposed. Approaches based on queuing theory, simulation models and theoretical distributions have been used in [30, 31, 15]. Time series models such as Autoregressive Integrated Moving Average (ARIMA) have been used since the late 1980's to forecast calls for emergency rescues, for example by Channouf et al. [32]. These models, however, come with restrictive data assumptions. Vile et al. have studied Singular Spectrum Analysis (SSA) [33]. They show that the approach outperforms traditional methods for strategic demand forecasts and can also be applied on the operational

planning level.

Baker and Fitzpatrick are one of only a few who try to predict emergency as well as non-emergency service demand using the Winter's exponential smoothing model [29].

A comprehensive mathematical model to predict the demand for emergency rescues was proposed by Ingolfsson [34]. It takes the expected fluctuations of response times and workload over time into account. Steins et al. have investigated the influence of different parameters like population, income, work places, road traffic or weather conditions for three Swedish EMS regions [1].

Zhou and Matteson proposed a fine-scale forecasting approach based on spatial-temporal kernel density estimation [35]. In order to predict demand over continuous time and discretised locations, Zhou et al. have employed a bivariate Gaussian mixture model [36]. A spatial-temporal Gaussian process regression is also described by Nabarro et al. [37].

Few papers have applied machine learning approaches for demand predictions. Setzler et al. have tested Artificial Neural Networks (ANNs) and found that they can be capable of producing accurate forecasts for small areas [25]. Zhou proposes three methods based on Gaussian mixture models, kernel density estimation, and kernel warping to provide spatio-temporal predictions with improved accuracy [38]. Chen and Lu and Chen et al. compare Moving Average, Artificial Neural Network, Linear Regression, and Support Vector Machine approaches for the forecast of emergency medical demand [39, 28]. They propose to use a Geographic Information System to manage and visualise the forecasts. Letzner uses machine learning approaches to predict to which hospital a patient is transported in case of a medical emergency [40].

Process times forecasting

Response times are a major performance indicator for EMS systems, with the travel time usually being the largest component of the response time [34]. It has been identified as one of main factors for the system quality. Henderson and Mason [41] as well as Harewood [42] based the prediction of travel times on the characteristics of road types and travel conditions encountered when travelling from the dispatch location to the scene of the incident. Fujiwara et al. and Aringhieri et al. simply use deterministic distances scaled by correction factors [43, 44].

Only very few authors have dealt with predictions for service times at an emergency scene, potentially

due to data limitations [22]. For example, Harper and Knight used Coxian phase-type distributions to study the effect of individual components of the emergency medical service process [45].

3. Taxonomy

Following Davenport [46], the key questions addressed by analytics for EMS services are displayed in Table 1. Reports of the past can tell for example, if the response time target was met together with the average and maximum response time. Additionally, the utilisation of ambulances is of interest. Analysing the driving times and dispatching decisions can then give an insight on the ambulance utilisation and the actual response times. Keeping track of the present should include the current traffic in order to predict driving times and determining shortest routes to incidents, the hospitals or back to the bases accordingly. In addition, an alert for above-average utilisation might be helpful in order to decide if additional resources will be required. Using optimisation approaches, recommendations on the next best action include the dispatching and relocation of ambulances depending on the current system status. Extrapolation is important for EMS logistics at all planning levels. The expected demand for emergency services is the main input parameter for models and algorithms that determine the number and locations of ambulances and bases. With a simulation of the whole EMS system different future scenarios together with varying location solutions, policies and (dispatching/relocation) strategies can be studied and tested before applying them in practice.

In the taxonomy, we now want to focus on predictive analytics, i.e. forecasting, in EMS logistics. We propose to distinguish the following three dimensions:

- Service: Emergency rescue or patient transport.
- Entity: Demand or process times (e.g. driving times, service times, chute times, turnaround times).
- Planning level: Strategic, tactical or operational.

Note that definitions for the different time intervals in the response process for an emergency rescue can be found in [2], for example.

In addition, we define several external factors that influence the choice of a forecasting approach and the expected output quality:

- Data: data quality, number and variety of input parameters, years of data and granularity.

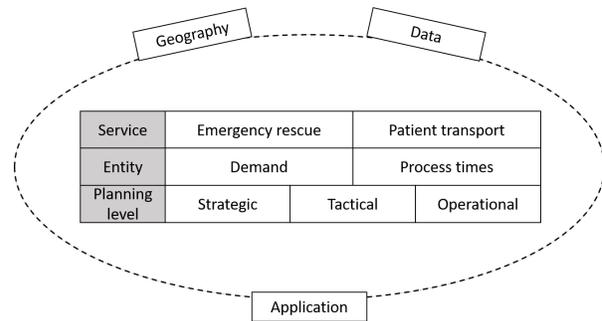


Figure 1. Taxonomy for EMS forecasting

- Application: is output used for a mathematical approach (and what are the characteristics) or displayed on a dashboard, for example.
- Geography: structure of the EMS region, size of the region, distribution of the population and demand throughout the area etc.

Ambulance location problems, for example, are strategic planning problems and often only use one or few input values for each demand point. These values represent the demand that is supposed to be covered within the given response time. These could be median values or the 95% quantile, for example. Often, also the driving times and service times are deterministic, while service times then represent the whole time interval from arrival at the scene to being available (back at the base) again. Other approaches like the stochastic program proposed by Nickel et al. [47] use scenarios for the strategic ambulance problem with each scenario representing an hour of a day, for example. Then, forecasts are necessary that determine probabilities that a specific number of ambulances are called to one demand point in one hour.

On the operational level, coordination centre managers and dispatchers would benefit from a dashboard showing the expected demand for emergency rescues in the next 30 to 60 minutes. Precise short-term forecasts are necessary that should make use of current external values like the time, day of the week, month, weather, traffic etc., but also of the demand that was observed in the past hours.

For the two described use cases different requirements for the forecasting methods and results exist. It is therefore reasonable to assume that different methods will lead to the best results.

Figure 1 summarises the taxonomy. Table 2 presents the most relevant use cases, i.e. planning problems as well as information for coordination centres.

Most authors aim at improving the forecast quality without explicitly stating the use case, i.e.

Table 1. Key questions addressed by analytics (based on [46])

	Past	Present	Future
Information	What happened?	What is happening now?	What will happen?
	(Reporting)	(Alerts)	(Extrapolation)
	<i>Compliance with the response time target; utilisation of the ambulances</i>	<i>Current traffic (traffic jams), alert for above-average utilisation</i>	<i>Expected demand</i>
Insight	How and why did it happen?	What is the next best action?	What is the best/worst that can happen?
	(Modeling, experimental design)	(Recommendation – optimisation)	(Prediction, optimisation, simulation)
	<i>Analysis on the driving times and the dispatching</i>	<i>Ambulance dispatching, relocation</i>	<i>Simulation of the whole EMS system</i>

Table 2. Planning problems and information sources in taxonomy structure

		Emergency services	Patient transports							
Entity	Planning level	Strategic	Tactical	Operational						
	Process times	Demand	Strategic reports	Ambulance location problem	Ambulance location problem	Real-time reports on Dashboard	Ambulance relocation problem	Ambulance dispatching	Patient transport problem	
Call centre and ambulance workforce				Staff rosters						
Service time										
Driving times										
Chute time										
Turnaround time										
Pick-up time										
Drop-off time					Patient transport problem					

planning problem or source of information, they are targeting. Spatial-temporal forecasting approaches with fine granularity are more detailed than the usual input for strategic planning problems, but there is a clear benefit at the operational level. So far, machine learning approaches have mainly been used for these types of forecast, e.g. [28, 38]. When time-series forecasts aim at predicting daily arrivals as in [32], the output can be applied to planning problems on the tactical level or accumulated to be used as input on the strategic level.

4. Initial results of strategic demand forecasting through machine learning

As introduced in Section 2.2, strategic demand forecasting has so far been mainly addressed through time series analysis and signal processing approaches. Given recent advances in the field of machine learning, however, we argue that machine learning may also be suitable for strategic demand forecasting. Consequently, in this section, we present initial results of our study comparing strategic demand forecasting through machine learning.

This study is performed in close collaboration with a German EMS coordination centre.

4.1. Problem definition

Results of strategic emergency demand forecasting approaches are mainly used as input for the strategic ambulance and base location problem. Ambulances and bases are located in such a way that the (expected) coverage is maximised and response times for emergency rescues are minimised. The demand to be covered is one of the main input parameters for these models. Often, the investigated EMS region is defined by a set of demand points representing sub-regions such as zip code areas. In the majority of the ambulance location models, one demand value per demand point is used as input. Therefore, applying strategic emergency demand forecasting approaches aims at determining this set of demand values, for example as median values or 95% quantiles derived from daily forecasts.

4.2. Data understanding

Data available for this study is based on historic calls from the collaborating EMS coordination center. The data reports the number of calls within time windows of one hour for each zip-code within the EMS region over a period of nearly six years (2013-01-01 to 2018-12-10). Given 89 zip-code areas, the data set consists of 4,635,120 data points that result in a total of 589,419 emergency calls and dispatched vehicles in the

EMS region in the given time period.

The overall call demand is depicted in Figure 2. We see a small increase in call volume over the past years with a higher increase at the end of 2018.

Aligned with the approach presented in [24], Figure 3 depicts the average weekly demand for each hour within the week (i.e. $24 \frac{h}{d} \times 7 \frac{d}{w} = 168 \frac{h}{w}$). On a visual note, we clearly see an increase in calls during daytime on weekdays and an increased call volume at Friday and Saturday night.

4.3. Forecasting models

We used three different forecasting techniques: First—as a time series based forecasting technique—we used the recently introduced *Prophet* approach [48]. In order to represent machine learning based approaches, we used, second and third, *Regression Trees* as well as *XGBoost Regression*.

The prediction pipeline is depicted in Figure 4: Whilst we trained individual prophet, Regression Tree, and XGBoost Regression models for each individual zip-code area, we also leveraged the machine learning based methods capability of multi-dimensional input and created—in addition to the zip-code individual models—one additional Regression Tree and XGBoost Regression model that receives all data (i.e. zip-code as an independent variable) as input. In the following, we refer to the zip-code individual Regression Tree and XGBoost Regression models as Regression Tree (I) and XGBoost Regression (I) and to the overall Regression Tree and XGBoost Regression model as Regression Tree (O) and XGBoost Regression (O) model. Instead of providing a holistic timestamp to machine learning based methods, we split up the timestamp into its individual components and provided them individually (i.e. year, month, day, hour) as well as the week of the year and the weekday.

4.4. Results

Table 3 shows initial results of our comparison with very simple machine learning methods. As performance measures, we use mean-squared-error (MSE), mean-average-error (MAE), as well as root-mean-squared-error (RMSE).

Overall, the results in Table 3 can be summarised as follows: First, all regression tree models (i.e. Regression Tree (I) as well as Regression Tree (O)) are greatly outperformed by the Prophet and XGBoost Regression models. Second, there is little performance difference between zip-code specific and overall demand prediction model (i.e. (I) vs. (O) models). Third, depending on performance measure, the

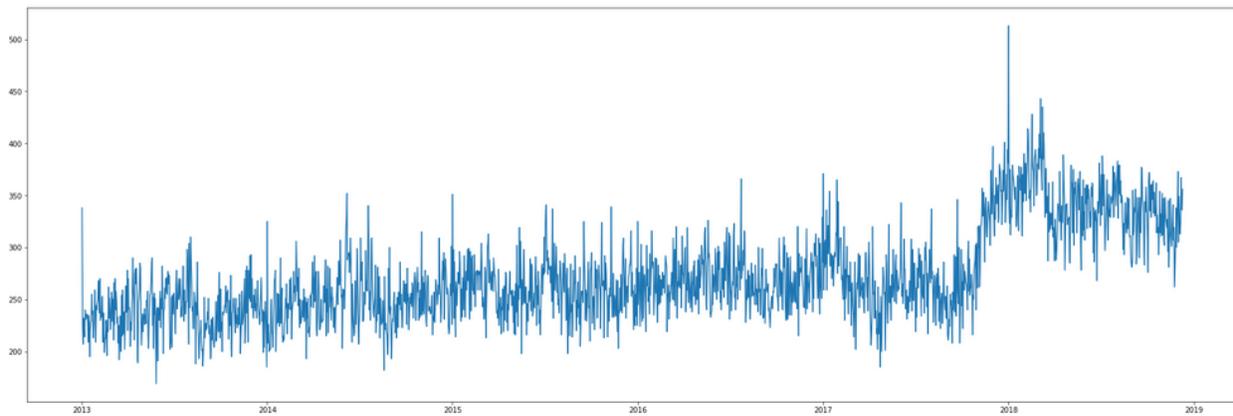


Figure 2. Overall Demand

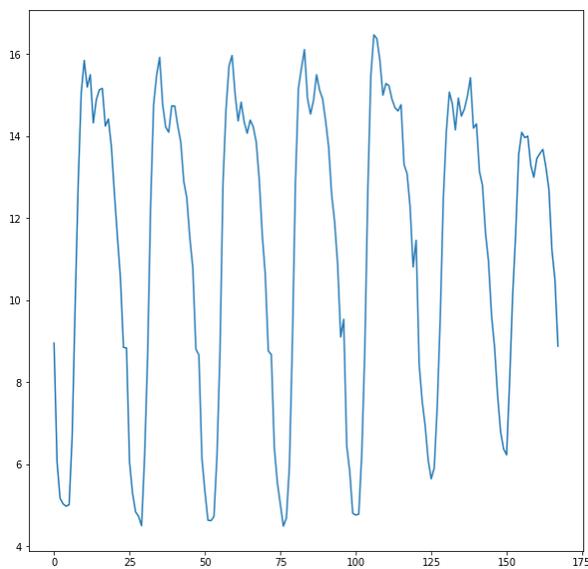


Figure 3. Average Weekly Demand

XGBoost models perform both, better as well as worse than the time-series based Prophet models.

4.5. Agenda for future research

As future research, a thorough analysis and comparison of forecasting approaches for all cases of the taxonomy is desirable. Our findings lead to the following research questions:

RQ1 Which forecasting method performs best for which use case?

RQ2 Can machine learning approaches outperform the conventional forecasting methods?

RQ3 Can we define necessary levels for data

availability and quality for the use cases in order to ensure the forecasting performance?

RQ4 Does including social media and telephone data improve the forecasting performance for emergency rescue demand on all planning levels?

In a first step, we plan on introducing additional data sources to enhance our strategic forecasting models for emergency demand. In addition, we will focus on operational demand forecasting and analyse the performance of machine learning approaches. While time series forecasts that performed well on the strategic level can only use a single input, for example the past demand over time, machine learning approaches can use multi-dimensional input. Another important aspect of the investigation is to analyse the usability of social media and telephone data for strategic as well as operational EMS forecast, as for example proposed by [49].

As a third use case we want to investigate approaches for operational patient transport forecasts. Due to several reasons, many transport tasks are not known in advance and have to be carried out on short notice in Germany. Therefore, scheduling the tasks is challenging and it is also difficult to determine how many ambulance will be necessary and if transports can be fulfilled within the desired time frame.

As output of RQ1 we aim at filling Table 2 with the most suitable approaches, taking also data characteristics potential applications into account.

5. Conclusion

EMS systems worldwide are very complex, but important for society's health and welfare. The application of adequate methods and algorithms to solve implied planning problems can improve the service

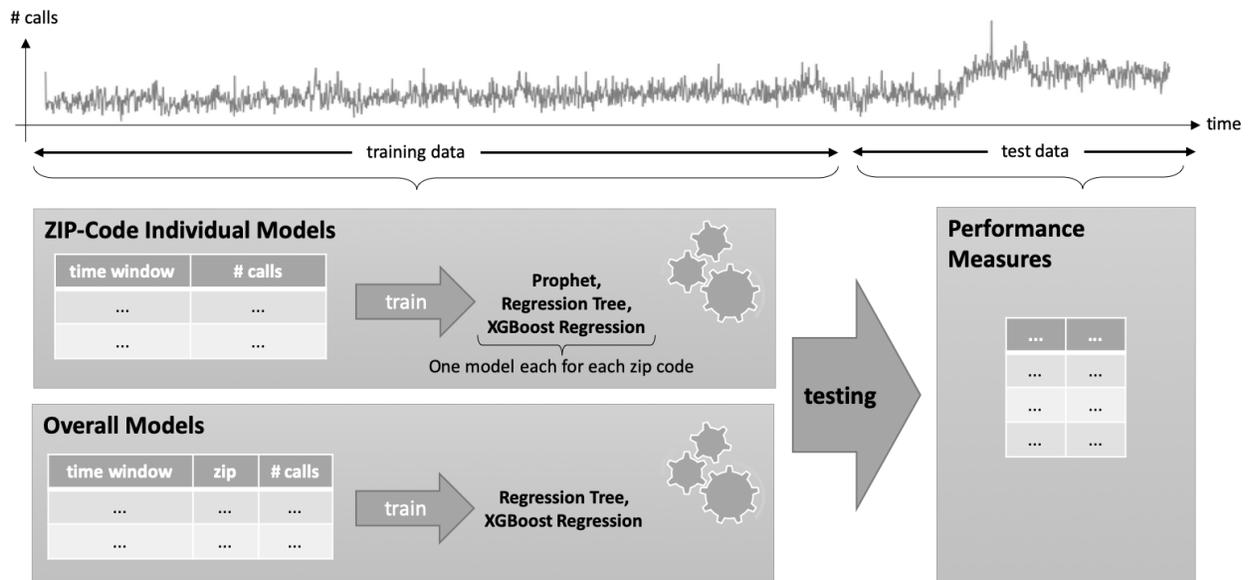


Figure 4. Forecasting Pipeline

Table 3. Initial Forecast Results

	Prophet	Regression Tree (I)	Regression Tree (O)	XGBoost Regression (I)	XGBoost Regression (O)
MSE	0.177847	0.333227	0.332337	0.184315	0.185867
MAE	0.225752	0.240712	0.332337	0.224038	0.226874
RMSE	0.421720	0.577258	0.576487	0.429320	0.431123

quality, i.e., reduce response / waiting times as well as installation cost. As solutions of optimisation methods strongly cohere with the input data, substantiated data analysis and forecasts are crucial for high quality solutions. Therefore, predictive and prescriptive analytics must be used together in order to improve the performance of EMS systems.

In this paper we have presented a taxonomy for EMS forecasting and commented on the existing literature accordingly. We have also performed a first analysis for strategic demand forecasting through machine learning with data from a German EMS region. Based on our findings, we have defined a set of research questions to be answered in future research.

As next steps we want to further investigate forecasting approaches for German EMS regions. We will test, analyse and compare different methods, including machine learning approaches. The aim is to define the most promising approaches for the different forecasting problems. We will also investigate the importance of data quality, granularity and variability.

References

- [1] K. Steins, N. Matinrad, and T. Granberg, "Forecasting the demand for emergency medical services," in *Proceedings of the 52nd Hawaii International Conference on System Sciences*, 2019.
- [2] M. Reuter-Oppermann, P. L. van den Berg, and J. L. Vile, "Logistics for Emergency Medical Service systems," *Health Systems*, vol. 6, no. 3, pp. 187–208, 2017.
- [3] C. Toregas, R. Swain, C. ReVelle, and L. Bergman, "The location of emergency service facilities," *Operations Research*, vol. 19, no. 6, pp. 1363–1373, 1971.
- [4] M. S. Daskin, "A maximum expected covering location model: Formulation, properties and heuristic solution," *Transportation Science*, vol. 17, no. 1, pp. 48–70, 1983.
- [5] E. Erkut, A. Ingolfsson, and G. Erdoğan, "Ambulance location for maximum survival," *Naval Research Logistics (NRL)*, vol. 55, no. 1, pp. 42–58, 2008.
- [6] M. Gendreau, G. Laporte, and F. Semet, "A dynamic model and parallel tabu search heuristic for real-time ambulance relocation," *Parallel Computing*, vol. 27, pp. 1641–1653, Nov. 2001.
- [7] C. Jagtenberg, S. Bhulai, and R. van der Mei, "An efficient heuristic for real-time ambulance redeployment," *Operations Research for Health Care*, vol. 4, pp. 27–35, Mar. 2015.
- [8] L. Brotcorne, G. Laporte, and F. Semet, "Ambulance location and relocation models," *European Journal of Operational Research*, vol. 147, no. 3, pp. 451–463, 2003.

- [9] G. Carter, J. Chaiken, E. Ignall, and N. Jun, "Response Areas for Two Emergency Units," *Operations Research*, vol. 20, no. 3, pp. 571–594, 1972.
- [10] S. Dean, "Why the Closest Ambulance Cannot be Dispatched in an Urban Emergency Medical Services System," *Prehospital and Disaster Medicine*, vol. 23, no. 02, pp. 161–165, 2008.
- [11] T. Andersson and P. Värbrand, "Decision support tools for ambulance dispatch and relocation," *Journal of the Operational Research Society*, vol. 58, pp. 195–201, Mar. 2007.
- [12] C. Jagtenberg, S. Bhulai, and R. van der Mei, "Dynamic ambulance dispatching: is the closest-idle policy always optimal?," *Health care management science*, pp. 1–15, 2016.
- [13] V. Bélanger, A. Ruiz, and P. Soriano, "Recent optimization models and trends in location, relocation, and dispatching of emergency medical vehicles," *European Journal of Operational Research*, vol. 272, no. 1, pp. 1–23, 2019.
- [14] P. Bradbeer, C. Findlay, and T. Fogarty, "An ambulance crew rostering system," in *Real-World Applications of Evolutionary Computing*, pp. 267–279, Springer, 2000.
- [15] H. K. Rajagopalan, C. Saydam, H. Setzler, and E. Sharer, "Ambulance deployment and shift scheduling: An integrated approach," *Journal of Service Science and Management*, vol. 4, no. 01, p. 66, 2011.
- [16] M. van Buuren, G. J. Kommer, R. van der Mei, and S. Bhulai, "Ems call centers with and without function differentiation: a comparison," *Operations Research in Health Care*, vol. 12, pp. 16–28, Mar. 2017.
- [17] S. Parragh, J.-F. Cordeau, K. Doerner, and R. Hartl, "Models and algorithms for the heterogeneous dial-a-ride problem with driver-related constraints," *OR Spectrum*, vol. 34, no. 3, pp. 593–633, 2012.
- [18] L. Ardekani, D. Haight, A. Ingolfsson, M. Salama, and M. Stanton, "Scheduling and routing ambulances that provide inter-facility patient transfers," tech. rep., 2014.
- [19] Y. Kergosien, C. Lente, D. Piton, and J.-C. Billaut, "A tabu search heuristic for the dynamic transportation of patients between care units," *European Journal of Operational Research*, vol. 214, no. 2, pp. 442–452, 2011.
- [20] M. Schilde, K. Doerner, and R. Hartl, "Metaheuristics for the dynamic stochastic dial-a-ride problem with expected return transports," *Computers & Operations Research*, vol. 38, no. 12, pp. 1719–1730, 2011.
- [21] Y. Kergosien, V. Bélanger, P. Soriano, M. Gendreau, and A. Ruiz, "A generic and flexible simulation-based analysis tool for ems management," *International Journal of Production Research*, vol. 53, no. 24, pp. 7299–7316, 2015.
- [22] L. Aboueljinane, E. Sahin, and Z. Jemai, "A review on simulation models applied to emergency medical service operations," *Computers and Industrial Engineering*, vol. 66, pp. 734–750, 2013.
- [23] N. Matinrad, *An Operations Research Approach for Daily Emergency Management*. PhD thesis, Linköping University Electronic Press, 2019.
- [24] D. Matteson, M. McLean, D. Woodard, and S. Henderson, "Forecasting emergency medical service call arrival rates," *The Annals of Applied Statistics*, vol. 5, no. 2B, pp. 1379–1406, 2011.
- [25] H. Setzler, S. Park, and C. Saydam, "EMS call volume predictions: A comparative study," *Computers & Operations Research*, vol. 36, pp. 1843–1851, 2009.
- [26] J. Vile, J. Gillard, P. Harper, and V. Knight, "Time-dependent stochastic methods for managing and scheduling emergency medical services," *Operations Research for Health Care*, 2015.
- [27] J. Lowthian, D. Jolley, A. Curtis, A. Currell, P. Cameron, J. Stoelwinder, and J. McNeil, "The challenges of population ageing: accelerating demand for emergency ambulance services by older patients, 1995-2015," *The Medical Journal of Australia*, vol. 194, no. 11, pp. 574–578, 2011.
- [28] A. Y. Chen, T.-Y. Lu, M. H.-M. Ma, and W.-Z. Sun, "Demand forecast using data analytics for the preallocation of ambulances," *IEEE journal of biomedical and health informatics*, vol. 20, no. 4, pp. 1178–1187, 2016.
- [29] J. Baker and K. Fitzpatrick, "Determination of an optimal forecast model for ambulance demand using goal programming," *Journal of Operational Research Society*, vol. 37, no. 11, pp. 1047–1059, 1986.
- [30] C. Bell and D. Allen, "Optimal planning of an emergency ambulance service," *Socio-Economic Planning Science*, vol. 3, pp. 95–101, 1969.
- [31] R. C. Larson, "A hypercube queuing model for facility location and redistricting in urban emergency services," *Computers and Operations Research*, vol. 1, no. 1, pp. 67–95, 1974.
- [32] N. Channouf, P. L'Ecuyer, A. Ingolfsson, and A. Avramidis, "The application of forecasting techniques to modelling Emergency Medical System calls in Calgary, Alberta," *Health Care Manage Science*, vol. 10, no. 1, pp. 25–45, 2007.
- [33] J. Vile, J. Gillard, P. Harper, and V. Knight, "Predicting ambulance demand using singular spectrum analysis," *Journal of the Operational Research Society*, vol. 63, no. 11, pp. 1556–1565, 2012.
- [34] A. Ingolfsson, *Operations Research and Health Care Policy*. EMS Planning and Management, Springer New York, 2013.
- [35] Z. Zhou and D. S. Matteson, "Predicting ambulance demand: A spatio-temporal kernel approach," in *Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining*, pp. 2297–2303, ACM, 2015.
- [36] Z. Zhou, D. S. Matteson, D. B. Woodard, S. G. Henderson, and A. C. Micheas, "A spatio-temporal point process model for ambulance demand," *Journal of the American Statistical Association*, vol. 110, no. 509, pp. 6–15, 2015.
- [37] S. Nabarro, T. Fletcher, and J. Shawe-Taylor, "Spatiotemporal prediction of ambulance demand using gaussian process regression," *arXiv preprint arXiv:1806.10873*, 2018.
- [38] Z. Zhou, "Predicting ambulance demand: Challenges and methods," *arXiv preprint arXiv:1606.05363*, 2016.
- [39] A. Y. Chen and T.-Y. Lu, "A gis-based demand forecast using machine learning for emergency medical services," *Computing in Civil and Building Engineering (2014)*, pp. 1634–1641, 2014.
- [40] J. Letzner, "Analysis of emergency medical transport datasets using machine learning," 2017.

- [41] S. G. Henderson and A. J. Mason, *Ambulance Service Planning: Simulation and Data Visualisation*, pp. 77–102. Springer US, 2004.
- [42] S. Harewood, “Emergency ambulance deployment in Barbados: A multi-objective approach,” *The Journal of the Operational Research Society*, vol. 53, no. 2, pp. 185–192, 2002.
- [43] O. Fujiwara, T. Makjamroen, and K. Gupta, “Ambulance deployment analysis: A case study of Bangkok,” *European Journal of Operational Research*, vol. 31, no. 1, pp. 9–18, 1987.
- [44] R. Aringhieri, G. Carello, and D. Morale, “Ambulance location through optimization and simulation: The case of Milano urban area,” in *The 38th annual conference of the Italian operations research society optimization and decision sciences*, Università degli Studi di Milano, Polo didattico e di ricerca di Crema, 2007.
- [45] P. Harper and V. Knight, “Modelling Emergency Medical Services with phase-type distributions,” *Health Systems*, vol. 1, pp. 58–68, 2012.
- [46] T. Davenport, J. Harris, and R. Morison, *Analytics at work: Smarter decisions, better results*. Harvard Business Press, 2010.
- [47] S. Nickel, M. Reuter-Oppermann, and F. Saldanha-da Gama, “Ambulance location under stochastic demand: A sampling approach,” *Operations Research for Health Care*, 2015.
- [48] S. J. Taylor and B. Letham, “Forecasting at scale,” *The American Statistician*, vol. 72, no. 1, pp. 37–45, 2018.
- [49] K. Eismann, M. Reuter-Oppermann, and K. Fischbach, “Towards Social Media Decision Support for Joined EMS and Crisis Logistics,” in *Logistics and Supply Chain Management in Crisis Response Proceedings of the 16th ISCRAM Conference* (J. J. G. Zeno Franco and J. H. Canós, eds.), 2019.