Metamodel Optimization of a Complex, Rural-Urban Emergency Medical Services System

Matthew Snyder[a], Byran J. Smucker[a]

[a] *Department of Statistics, Miami University, Oxford, Ohio,*

Abstract

Complex simulation systems, such as those involving emergency medical services (EMS), are often too computationally demanding to be used in optimization problems. Metamodeling is an attractive alternative, in which a sample of system configurations is evaluated using simulation, and a fast predictive model is developed as a surrogate for the slow simulator. Though the metamodeling literature is extensive, there has been little exploration of how much data is required to construct metamodels that can be used to solve optimization problems effectively, particularly in the context of a complicated rural-urban EMS system environment. In this work, the EMS system in northern St. Louis County, Minnesota, is studied, with the goal of discovering station configurations with improved response times. The underlying physical system is complex, with 12 stations spread across both rural and urban areas and a fairly large geographic footprint. A decade of call data is used to develop and validate a stochastic discrete event simulator (DES) for this system, and then the simulator and raw data is used to select realistic station configurations to train the metamodel. Results are first given for just a single station within the system, and then increasingly complex settings are examined culminating with consideration of all 12 stations. Overall, though the metamodeling approach was effective for simpler cases, it requires a tremendous amount of data for larger settings. Specifically for the St. Louis County example, improved configurations were found for the one- and two-station cases, but the amount of data required to produce effective metamodels for the five- and twelve-station versions of the system was computationally infeasible given current DES and optimization heuristic implementations.

Keywords: Emergency medical services, Discrete event simulation, Metamodeling, Particle swarm optimization, Random forest

1. Introduction

An Emergency Medical Services (EMS) department is tasked with responding quickly to the medical needs of people in its community. Especially in life-threatening emergencies, faster response times can translate to higher chances of survival, which makes reducing this time a priority for EMS systems [1, 2]. In particular, organizations and governing agencies, such as the National Fire Protection Association (NFPA), publish guidelines that specify 90th percentile response time thresholds [3, 4]. These response times depend on a system's available resources, including ambulances and staff, as well as the locations of the stations with respect to the calls [5, 6]. It has been shown that adjusting the station locations or available equipment from a finite list of candidate configurations can reduce response times and save lives [2, 7, 8, 4].

Many types of models and simulations have been developed to test alternative resource location, allocation, and dispatching policies [9, 10, 11]. While mathematical models are fast and can be used to solve optimization problems directly, they are based on many simplifying assumptions. Simulations, including discrete event simulation (DES), are typically much more flexible and better able to approximate such complex and interconnected systems, but are more computationally intensive which becomes a limiting issue for large-scale systems. In fact, the computational difficulty entailed in this problem, particularly as encountered in a commercial application, is what motivated the current research. Indeed, attempts to use a DES to directly optimize station configurations under simplified settings were unsuccessful due to memory and connection limitations associated with the large number of simulation runs.

Thus, this paper seeks to develop a simplified predictive model, called a metamodel, which only considers the relationship between the basic inputs and outputs of the simulator, rather than accounting for each of the detailed, complex aspects of the simulation. As such, metamodels can be evaluated much more quickly [12] than the DES. The main contribution of the current research, motivated by extensive work on a real EMS system in northern St. Louis County, MN, USA, as well as by a company doing analytics and optimization work for EMS systems, is to study the amount of training data needed to fit a metamodel well enough to usefully optimize station locations. As discussed in Section 2, there is little in the literature addressing the nexus of issues confronted here: optimization using a metamodel built to mimic a DES, which in turn is modeling an underlying complex, rural-urban EMS system. Investigating the amount of data necessary to effectively construct such a metamodel has not been investigated in an EMS simulation setting, and rarely in the metamodeling literature more broadly.

The following section presents a review of related simulation and metamodeling work, Section 3 describes the methodology of the present study, and Section 4 contains the results, followed in Section 5 by a discussion of the findings, areas for future work, and a conclusion.

2. Related Work

As mentioned above, many types of models have been developed in the study of EMS systems, with simulation models more flexible and accurate than mathematical models. In particular, discrete event simulation (DES) has been shown to accurately model EMS [1, 13, 14, 15] and other healthcare systems [16, 17, 18, 19, 20]. DES is a method of modeling complex systems comprised of events in time, in which the system remains unchanged between events. It is relatively adaptable and can handle constrained resources, and it has been the most common simulation tool used to model EMS systems [11, 21], in comparison to continuous or constant time-step simulations [22]. Agent-based simulation (ABS), or a combination of ABS and DES, has also been used to model EMS systems as it can track and model the movement of ambulances in the system in more realistic and complex ways [22, 23, 24, 25] than DES alone.

These simulators can then be used to compare how various system configurations affect response times for an EMS system [26, 27, 28, 29, 30]. In the review by Li et al. [10], the authors include an overview of simulation techniques used to test different policies for the allocation and deployment of ambulances and compare multiple pre-defined ambulance location configurations, but do not consider a full optimization of station locations. Aringhieri et al. [11] provides a review of EMS ambulance location, relocation, and dispatching policy problems, including the use of simulation models and specifically DES. Of particular interest to the present work, the study by Mason [31] explores a simulation-optimization of vehicle base locations, but only performs a local optimization by perturbing existing station locations. Additional research has been performed specifically for EMS systems in heterogeneous rural-urban regions. Rather than treating all areas equivalently and only considering response time, models have been developed that vary response time and survival targets in different areas [32, 33, 34, 35], or employ additional optimization objectives such as maximizing coverage or survival [36, 37] in order to balance efficiency and equity throughout the system. As seen in these papers, although simulation allows for more flexibility and accuracy in modeling EMS systems, especially complex rural-urban systems, the added computational expense has limited the optimization of station locations to simply selecting from pre-defined candidate locations or implementing a local optimization procedure.

The use of metamodels as a surrogate for the simulation model can simplify the underlying structure while hopefully maintaining high accuracy to enable a full optimization. To be more specific, a metamodel is a predictive model that is trained by runs of the simulator at a diverse set of system configurations. Then, if enough information is provided by the simulations, the metamodel effectively mimics the much more complex and computationally difficult simulator and can be used for decision-making and even optimization. In recent years, metamodeling has been widely studied, and for EMS systems has been used to compare and optimize both response times and survival rates by changing ambulance locations [38] and dispatch policies [39]; these models were based on agent-based simulators. More generally, comparisons of different types of metamodels, such as regression splines, kriging, artificial neural networks, and random forests, as well as different sampling approaches, including space-filling designs and adaptive sampling, have been conducted in many contexts on a variety of simulators [40, 41, 42, 43, 44, 45], including discrete event simulators [46, 47]. Depending on the context of the problem and the ultimate goal, the recommended model and approach varies. In a comparison of several metamodeling strategies, for discrete event simulators of EMS systems, Hopkins and Smucker [48] found that k-nearest neighbors and random forest models yield the highest accuracy and predictive power.

The goal in metamodeling is to create a highly accurate model using the least amount of training data, since this training data is expensive to obtain. Jin et al. [49] was an early researcher of metamodel performance using different sample sizes and problem complexities. For large scale problems with at least ten predictors, the authors compare the accuracy of three sample sizes – scarce, small, and large – using R^2, relative average absolute error (RAAE), and relative maximum absolute error (RMAE). The samples are generated using Latin hypercubes, with sizes determined as functions of the number of parameters p, ranging from $3p$ to $3(p+1)(p+2)/2$. They found that the accuracy tended to increase as sample size increased, with average R^2 values near 70 for the large scale, nonlinear problems. Yang et al. [50] compared the RMSE of five types of metamodels fit using training sets with sizes ranging from $3p$ to $36p$, where $p = 4$, for a complex, nonlinear finite element model. Similarly, Kim et al. [51] studied the accuracy of metamodels built on samples of $3p$, $5p$, and $7p$ points using RMSE for sample mathematical problems with two through eight predictor variables. Kianifar and Campean [52] provided a recent, comprehensive literature review of metamodeling techniques with the goal of creating a guide for engineering professionals; the authors compared many facets of metamodeling for several mathematical-based engineering problems, including two different samples sizes (10- and 30-times the number of predictors), using a normalized RMSE. The general consensus for these studies was that increasing the sample size resulted in higher accuracies, but at varying rates for different model types, problem complexities, and error types. Other studies have compared the accuracy of simulation-based metamodels built on several pre-defined sample sizes, typically ranging from several dozen to several hundred training points, and occasionally reaching a thousand points [53, 54, 55, 56, 57, 58]. One exception is Ding and Zhang [59], who explored large-scale simulation metamodeling in settings with 10, 20, 50, and 16,675 predictors, and tested sample sizes ranging from 200 points to 30,000 points. They compared the RMSE at increasing sample sizes using multiple sampling designs, and generally found sharp decreases in errors that eventually plateaued; however, the authors focused on metamodeling for simulation prediction, rather than optimization where alternative methods may be more efficient.

Once a metamodel is constructed, it can be used to solve optimization problems, as in Osorio and Chong [54] who optimize signal plans in simulation-based transportation systems, and Ju et al. [56], who optimize turbomachinery designs in Monte Carlo simulation problems. Zeinali et al. [60] created a metamodel to approximate a DES for emergency departments with the goal of minimizing patients' waiting times, and discovered a resource configuration that reduced waiting times by 48% for a particular hospital. The metamodel was fit using fewer than 100 points, and the near-optimal solution was found quickly.

Unlike the existing literature, the current research seeks to optimize the location of the EMS stations over a continuous region rather than over a finite list of candidate locations or a perturbation of the existing system. To accomplish this, metamodels are fit over the underlying simulation, using as input only the locations of the EMS stations. The question considered is: How much training data is required to obtain metamodels reliable enough for use in optimization? Due to the intrinsic, interconnected nature of the rural-urban EMS system considered, the problem is significantly more complicated than the studies in the current literature and requires substantially more data. Further, because traditional, commonly-used space-filling designs are not practical in this setting since they waste computational resources on illogical station configurations, alternative sampling methods must be considered.

3. Data and Methods

In order to evaluate the efficacy of this metamodeling approach and determine the amount of data needed to obtain useful results, consider four versions of a case study, each of increasing complexity. These include a simple setting with one station, a two-station setting that accounts for the interaction of stations, and a full twelve-station setting to model at least a simplified version of the entire EMS system of northern St. Louis County, Minnesota (results for an intermediate five-station scenario are also reported in the Supplementary Material). For each situation, a DES was created and a number of samples of station configurations were generated to run through the DES. This resulted in a simulated 90th percentile for each configuration that was then used as the response for the fitted metamodel. Once fit, the metamodel and particle swarm optimization were used to generate a set of potentially optimal station configurations, which were then validated on the DES. If improvements in the metamodel and proposed configurations could be expected, the process repeated with a larger sample size, as illustrated by Figure 1.

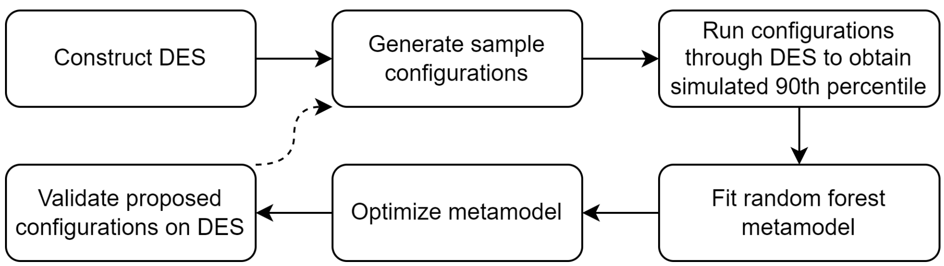

Figure 1: Flowchart of the methodology used in this article.

All data cleaning, simulation, model fitting, and optimization was done in R [61]. The packages `tidyverse` [62], `simmer` [63], `osrm` [64], `caret` [65], `randomForest` [66], and `pso` [67] were used extensively.

3.1. Data Handling

The data for this study was provided by Allen Lewis, Fire Chief and Emergency Manager of the Virginia Fire Department in Virginia, Minnesota. Calls related to medical emergencies between 2009 and 2019 were included, extracted from two sets of data. Calls collected between December 27, 2009 and April 3, 2016 were reported weekly, while those collected between January 1, 2018 and December 30, 2019 were reported daily. For every call, the date (or week), type of emergency, radio name of the responding vehicle, address, latitude, and longitude were recorded, along with the time that the call arrived as well as the times that the vehicle was dispatched, enroute, on scene, and cleared from the emergency.

Several additional variables were constructed based on the data. Both the station and vehicle type of the responding vehicle were extracted from the radio name; ambulances, Medical Response vehicles, and Battalion Chief vehicles were classified as "EMS", while other vehicles were recorded as "Fire". A binary city variable marked calls that were located in the cities of Virginia, Mountain Iron, Hibbing, Ely, Eveleth, or Chisholm and within 5 miles (8.1 km) of the station as true, while others were marked false. The time of day was recorded as "early morning" for calls arriving before 6:00 a.m., "morning" for calls arriving before noon, "afternoon" for calls arriving before 6:00 p.m., and "evening" for other times. The season was recorded as "winter" for calls arriving from December to February, "spring" for those arriving from March to May, "summer" for those arriving from June to August, and "fall" for those arriving from September to November.

To obtain accurate and reliable data for this study, the data was cleaned using the following procedures. In all, three datasets were constructed from the raw data. The first dataset, denoted A, omits calls with missing locations and missing or illogical times. It also omits calls whose times were judged to be unreasonable, when compared to the OpenStreetMap time. Dataset A is the most filtered set of calls (Table 1) and is used whenever it is necessary to have a set of calls with reliable times. More details regarding Dataset A are provided in Section 1 of the Supplementary Material. Another dataset, denoted B, was generated to assess the call frequency over time by computing the difference in call arrival times between unique emergencies. Only the 2018 and 2019 call data was used, as earlier data was reported weekly and not daily. Time of day and season variables were also computed as described above based on the arrival times. Finally, dataset C consists of call locations from the entire decade of call data. A binary city variable marked calls that were located in any city (Virginia, Mountain Iron, Hibbing, Ely, Eveleth, or Chisholm) as true, while others were marked false.

Together, these three datasets served as the foundation for the simulation and analysis of each version of the case study. Based on the number of stations in each scenario, datasets were filtered to include a subset of the responding stations. The one-station setting only accounted for the Virginia station; the two-station version added the Eveleth station; and the twelve-station setting encompassed the full system which includes the Aurora, Babbitt, Bois Forte, Buhl, Chisholm, Cook, Ely, Eveleth, Hibbing, Orr, Tower, and Virginia stations. (A five-station scenario is included in the Supplementary Material which includes Virginia and Eveleth as well as the Hibbing, Buhl, and Chisholm stations.) The two largest stations, Virginia and Hibbing, had four ambulances, while all other stations had two ambulances. Table 1 shows the amount of data in each dataset for each scenario based on this data handling. These sets of calls, locations, and available ambulances were used as the basis for the discrete event simulator described next.

Scenario	**Dataset A**	**Dataset B**	**Dataset C**	
	(Calls)	**(Calls)**	**(Calls)**	**(Unique Locations)**
One Station	6,240	8,788	22,335	5,285
Two Station	7,129	10,333	27,877	6,461
Five Station	16,524	18,768	49,640	11,588
Twelve Station	19,625	23,233	64,847	15,956

Table 1: Number of calls used for each dataset in each case study. Dataset A is filtered in order to have calls with reliable response times. Dataset B is a set of calls used to measure frequency and time between calls. Dataset C is a complete set of call locations.

3.2. Discrete Event Simulation

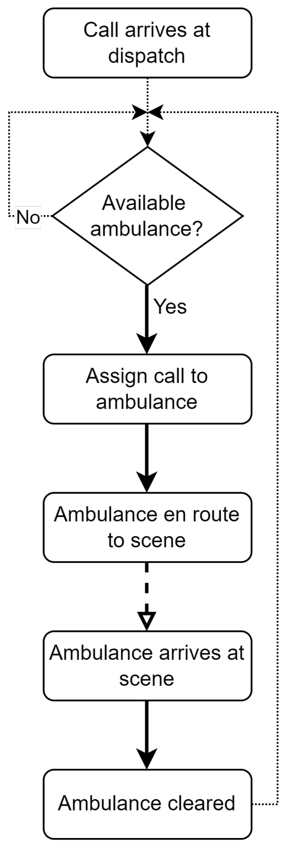

Figure 2: Discrete Event Simulation process. The various aspects of the response time are denoted by solid arrows. The travel time is denoted by the dashed arrow.

As described in the previous section, discrete event simulation (DES) is common in the EMS space. In this work, the use of DES is motivated by its appropriateness to the problem as well as collaboration with Levrum Data Technologies, a fire and EMS analytics company that employs similar tools. Thus, to model the EMS system, a DES was built using the `simmer` package [63] which corresponds to the first step in Figure 1. In an EMS system, the specific events modeled by the DES include the call arriving to the dispatcher, the dispatcher assigning an ambulance, the ambulance leaving for the scene, arriving at the scene, and being cleared from the emergency, as seen in Figure 2. The coordinates of the stations are the input to the DES, and the ambulances are the constrained resources. For this analysis, stations were assigned either two or four ambulances based on the current resources in Northern St. Louis County. Once a call arrives and is assigned to a station, an ambulance is seized until the call is cleared. If no ambulances are available, the call is queued and waits until a vehicle is cleared for the one-station case, or is assigned to the next closest station for the more complex cases. The location of the call is randomly selected from a list of all locations found in dataset C.

The time between call arrivals, dispatch to assignment times, assignment to enroute times, and onscene to clear times are randomly generated from gamma distributions and are represented by the solid arrows in Figure 2. These distributions were investigated in case they varied as a function of variables like time of day and time of year. The distribution for time between call arrivals was found to be based on time of day and season; assignment to enroute times based on time of day and vehicle type; and the dispatch to enroute and on scene to clear times based upon the binary city variable and vehicle type. As an example, the shape and rate parameter values for the gamma distributions used in the one-station setting are provided in Table 2. The remaining time needed to build the DES is the travel time, represented by the dashed arrow in Figure 2. This was estimated using a linear regression model that calibrated the OpenStreetMap times to realistic emergency response travel times, while accounting for variables like time of day and season. Based on a prediction interval from this regression analysis, a distribution was constructed from which the travel time was drawn. Further details are provided in the Supplementary Material, Section 2.

Putting all these fitted distributions together, the DES was constructed and run for one simulated year. The response time for each generated call was calculated as the difference in time between the call being assigned to a station and the ambulance arriving at the scene. Then, for a given configuration, the 90th percentile of the response times was computed across all calls since this is the primary metric of interest. In Section 3 of the Supplementary Material, the simulated 90th percentiles are compared to the 90th percentiles based on the historical data. Overall, the simulation provides response time distributions that are very similar to those found in the data.

Time Difference	**Factor Levels**	**Early Morning**		**Morning**		**Afternoon**		**Evening**	
		Shape	**Rate**	**Shape**	**Rate**	**Shape**	**Rate**	**Shape**	**Rate**
Call Arrival	**Spring**	1.271	0.007	0.768	0.005	0.889	0.010	1.106	0.011
	Summer	1.501	0.008	0.842	0.005	1.043	0.012	1.180	0.012
	Fall	1.234	0.007	0.745	0.005	0.972	0.011	1.208	0.013
	Winter	1.428	0.008	0.774	0.005	1.055	0.013	1.243	0.013
Assignment to Enroute	**EMS Vehicle**	1.572	0.817	0.824	0.468	0.952	0.868	1.191	0.943

		In City		**Not in City**	
		Shape	**Rate**	**Shape**	**Rate**
Dispatch to Enroute	**EMS Vehicle**	2.578	1.108	1.992	0.677
Onscene to Clear	**EMS Vehicle**	1.789	0.031	1.617	0.022

Table 2: Parameter values for the gamma distributions used in the discrete-event simulator at each factor level combination.

3.3. Metamodeling

Once constructed, the DES was used to generate a set of training data for the metamodeling process. This data came from a set of sample configurations that were run through the DES in order to obtain the simulated 90th percentiles, as illustrated in the second and third steps of Figure 1. Using the latitudes and longitudes of the station locations as inputs and the simulated 90th percentiles as outputs, the random forest metamodels were fit. Since the DES takes time to run, the goal is to fit the metamodels on the smallest set of training data that still yields accurate and informative results. Table 3 shows each of the training data sample sizes considered for the four settings.

Scenario	**Sample Sizes**
One Station	50, 100, 200, 300, 400, 500, 1000
Two Station	500, 1000, 2000, 3000, 4000, 5000
Five Station	2000, 10000, 50000
Twelve Station	5000, 10000, 50000

Table 3: Sample sizes studied for each version of the case study. Note that the five-station case is treated in the Supplementary Material.

Intuitively, a well-situated configuration that reduces the 90th percentile of response times would have stations located in the vicinity of the majority of calls, yet still spread throughout the region. Thus, rather than using a space-filling design to generate sample locations, a weighted sampling technique was used based on the spatial density of calls. This allows for more sample data in the intuitive locations, increasing the precision of the results in these areas. The spatial density of the calls were computed from dataset C using the `density.ppp` function of the `spatstat` R package [68], which increased the probability of choosing station locations in areas that have many emergency calls and decreased the probability of choosing irrelevant station locations. Details of the density estimation are provided in Section 4 of the Supplementary Material.

The one- and two-station versions simply selected one or two locations from this density for each configuration to generate the samples. For the twelve-station case (as well as the five-station scenario in the Supplementary Material), an additional constraint was added that forced all locations in a configuration to be at least four miles (6.4 km) apart, which is the distance between the two closest stations in the current configuration. This constraint was implemented due to the extremely high density of calls in the city of Virginia, which resulted in many configurations with several stations in very close proximity. While this added constraint introduces some limitations in the fitted metamodels, it allows for more exploration of intuitive configurations—those with stations spread throughout the region—while requiring less data. A brief exploratory study comparing several alternative sampling techniques is provided in Section 5 of the Supplementary Material, which found that using the constrained weighted technique ultimately resulted in the most promising proposed station configurations.

Once the sample of input configurations was generated, a random forest model was fit as the metamodel for each of the sample sizes considered, corresponding to the fourth step of Figure 1. Random forests were chosen due to their high predictive power in EMS settings [48]. Details on the implementation of the random forest are provided in Section 6 of the Supplementary Material.

3.4. Optimization

Optimizing the locations of the stations using the metamodel was attempted, via particle swarm optimization (specifically SPSO 2007), in order to find the station configuration that the metamodel predicts will have the shortest 90th percentile response time. Corresponding to the fifth step in Figure 1, this technique is based on a swarm of several configurations, called particles. The metamodel is evaluated for each particle, and the latitude and longitude of each station in the particle are assigned a velocity. The particles then move around the search space based on these velocities and are reevaluated by the metamodel, continuing until the global optimum is found. In particular, the optimization problem can be specified as follows:

$$\begin{aligned} &\text{Minimize } \; f(lat_1, long_1, lat_2, long_2, \ldots, lat_K, long_K) \\ &\qquad s.t. \;\; (lat_i, long_i) \in (\min_{n\in\mathcal{C}} y_n, \max_{n\in\mathcal{C}} y_n) \times (\min_{n\in\mathcal{C}} x_n, \max_{n\in\mathcal{C}} x_n),\; i = 1, 2, \ldots, K \end{aligned} \tag{1}$$

where f is the metamodel-predicted 90th percentile of the response times for a configuration with K stations, $K = 1, 2, 5, 12$; station k is located at latitude lat_k and longitude $long_k$; x_n and y_n are the longitude and latitude of call n in dataset C; and $\mathcal{C}$ is the set of all row indices in dataset C. Additional details of the PSO specification can be found in Section 7 of the Supplementary Material. Note that while the data used to fit the models included a distance constraint between stations for the five- and twelve-station cases, the optimization procedure did not, allowing stations to become close or even overlap.

Given the highly nonlinear and black-box nature of f, optimization in this setting is limited to the use of heuristics. Techniques including genetic algorithms [2, 69], tabu search [6], ant colony [10], and PSO [70], have been used to optimize models relating to EMS station locations and ambulance deployment, with GA and tabu search as the most popular methods [10]. In more general applications, comparisons between heuristics have often found PSO or GA to be the most promising techniques [21, 71, 72, 73], with PSO shown to more consistently find global optimal solutions. Thus, PSO is a reasonable approach to use in this setting.

For this study, the heuristic was run for each metamodel using 100 different initializations, with each procedure terminating after 50 iterations passed without improvement in the objective function. This technique provides a set of up to 100 different station configurations that, ideally, have low metamodel-predicted 90th percentile response times. Rather than proposing a single solution, it provides a set of configurations allowing for additional exploration and visualization of trends, which is desired since the use of metamodels has been shown to find improved, but not necessarily optimal, solutions for the underlying simulation system.

3.5. Validation

As the final step in Figure 1, each location proposed by the 100 metamodel optimizations of station locations (Section 3.4) was tested on the DES and compared to the current configuration. Due to the inherent randomness in the simulator, the configurations were run several times and the average 90th percentile response time was calculated. To start, the DES was run 100 times for the current station location and the mean and standard deviation of the 90th percentile response times were computed. Next, a simple statistical power analysis was conducted to determine how many runs of the DES were needed in order to detect differences of at least 30 seconds compared to the mean 90th percentile of the current station location, with 90% power and a significance level of 0.05. The DES was then run for each of the station configurations proposed by the optimization procedure, and average 90th percentile response times and corresponding 95% confidence intervals were compared to those from the current system.

Note that since the metamodels for the twelve-station case were substantially more complex than for the simpler cases, a majority of the 100 optimization procedures resulted in different configurations. Further, these simulations took additional time to run, so each proposed configuration was tested on the DES only once to start. Then, based on this single run, the configurations with a simulated 90th percentile below the upper 95% confidence interval bound for the current configuration were run on the DES additional times using the findings from the sample size analysis. Finally, the means and 95% confidence intervals of the simulated 90th percentile of response times were computed for these proposed configurations and compared to the current system.

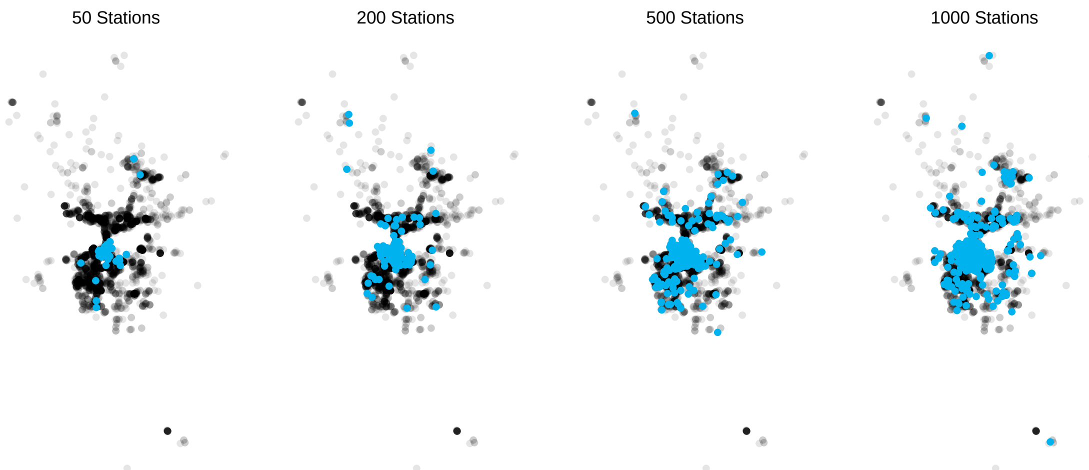

Figure 3: Sample locations (in blue) used as inputs for the one-station setting metamodels, overlaid on the historic call locations (in black).

4. Results

Using the methodology described in Section 3, this section reports on the results for the study of metamodeling as a strategy to perform optimization in a complex simulation environment. Included are the one-, two-, and twelve-station scenarios (with a five-station version in the Supplementary Material) so that metamodeling data needs can be better understood in increasingly complex environments.

4.1. One Station

In order to fit a random forest metamodel over the DES, weighted samples of seven incremental sizes were generated as inputs, as described in Section 3.3. Figure 3 shows the samples of size 50, 200, 500, and 1,000 in blue on top of all possible call locations shown in black. For samples of 200 and fewer locations, the stations were primarily located in the center of the region, while increasing the sample sizes to 500 or 1,000 expanded the coverage to the outlying areas as well. The DES was then run at each location, each run resulting in a corresponding 90th percentile response time. The metamodel, then, was fitted to each simulated dataset using the station locations as inputs and the 90th percentile response times as outputs. Figure 4a shows the goodness of fit of the models by sample size using the cross-validated R^2 of the selected random forest metamodel along with standard error bars. Although the R^2 is relatively high for a sample size of 50, it drops to a mean of 64% for a sample size of 100. It then approaches 80% for samples of size 200 and 300 before peaking with values around 88% for the larger sample sizes. The variability in the R^2 values is smallest for samples of size 400 and above.

Next, the station locations were optimized using each of the seven metamodels, the particle swarm optimization procedure, and 100 different initializations, which resulted in a set of proposed station locations. Each of these locations with a unique 90th percentile based on the metamodel, as well as the current location,

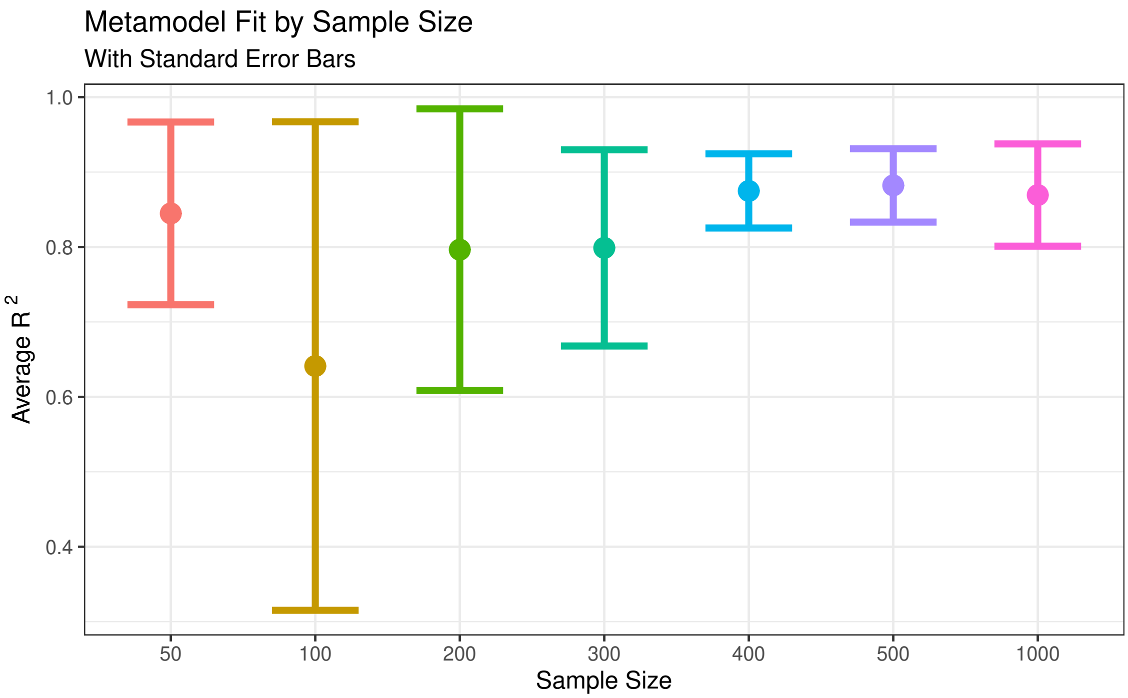

(a)

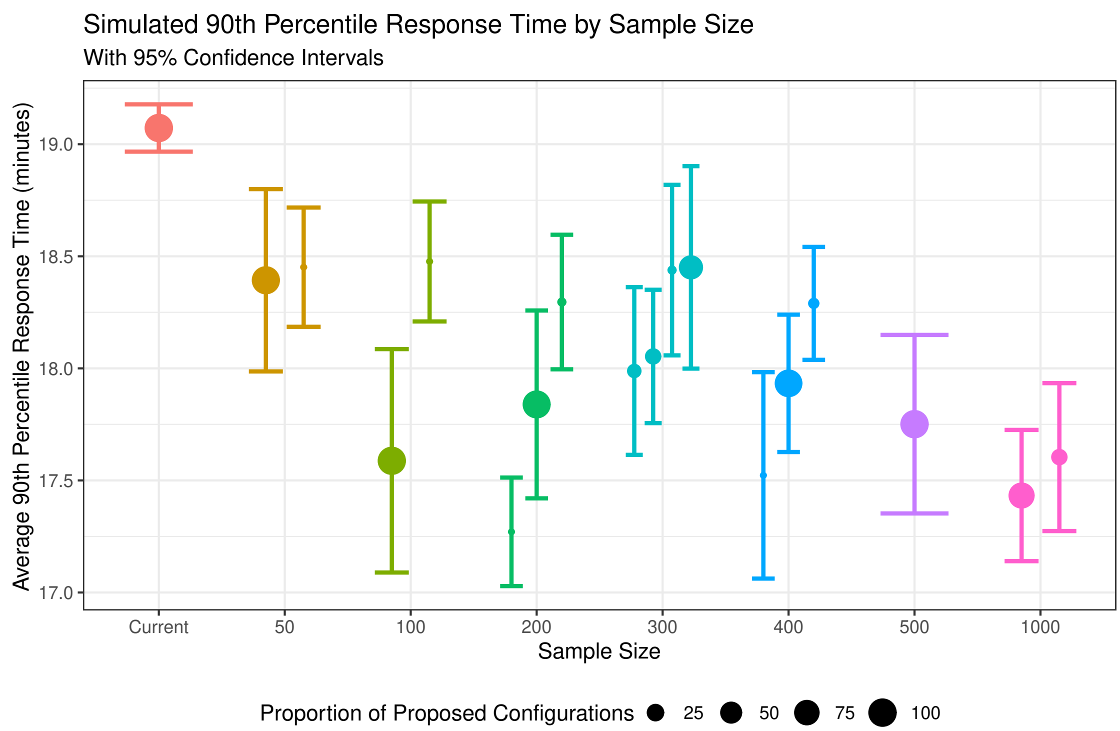

(b)

Figure 4: For the one-station setting, (a) metamodel fit by sample size with standard error bars; and (b) average DES simulated 90th percentile response times for proposed locations.

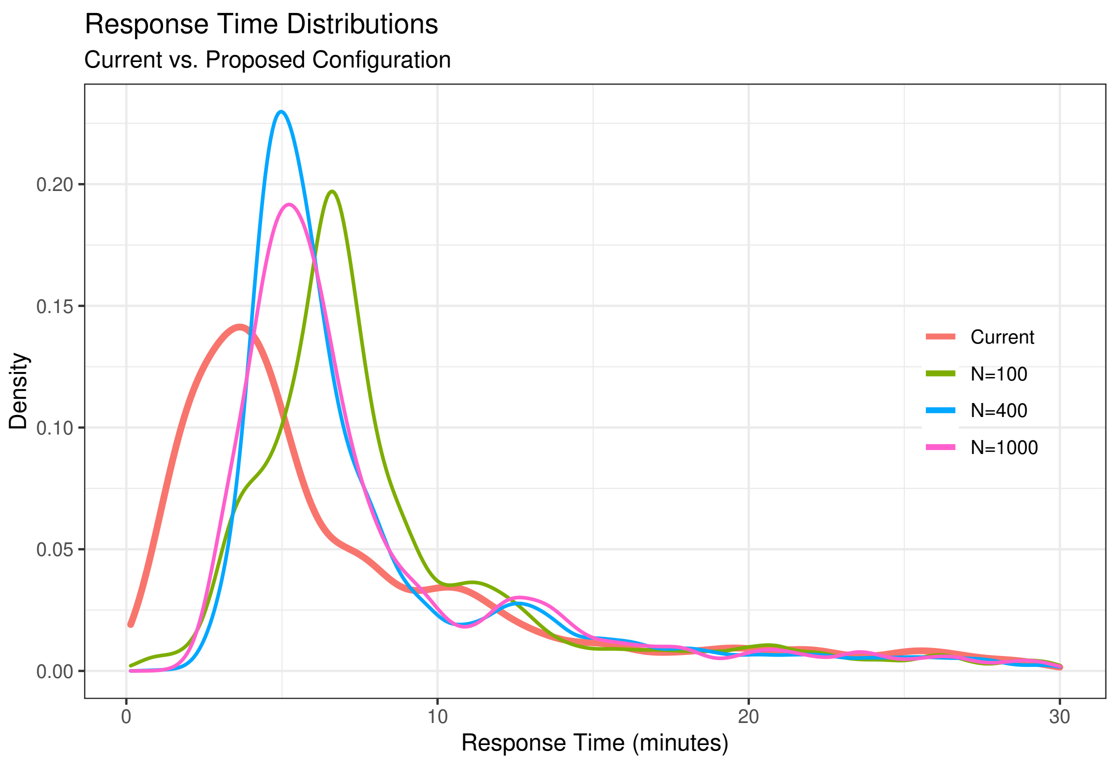

(a)

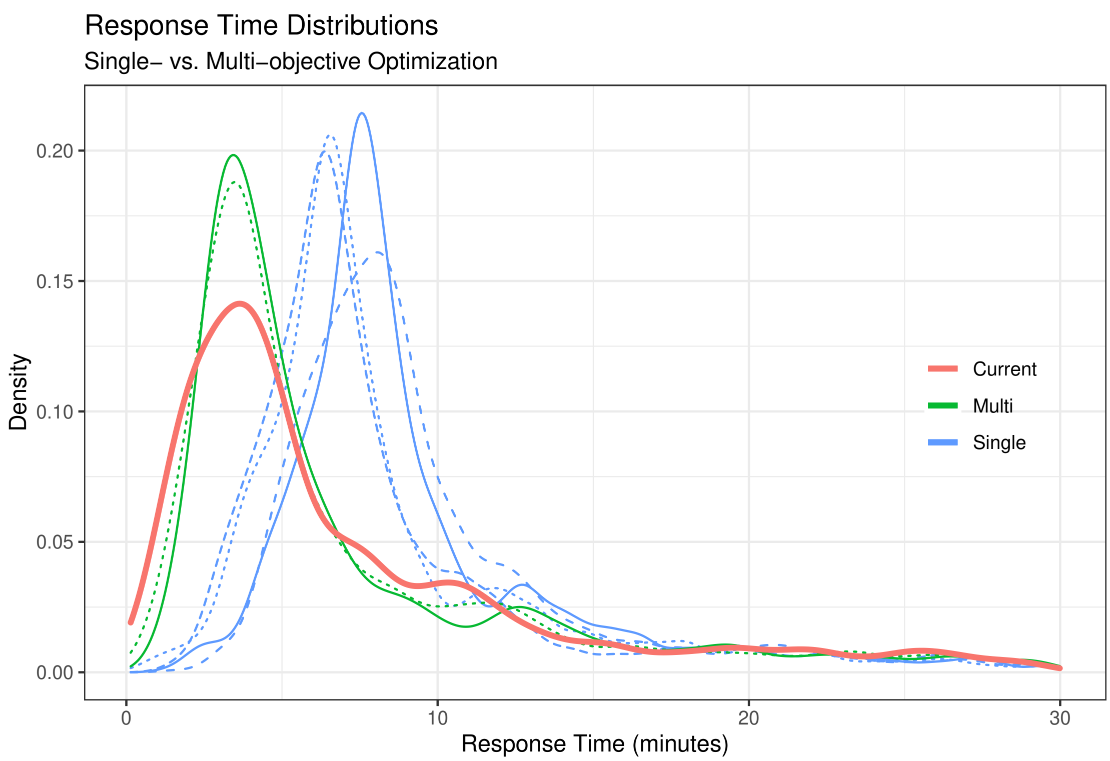

(b)

Figure 5: For the one-station setting, (a) distributions of response times for the current and proposed configurations; and (b) distributions of response times for proposed configurations based on the single- and multi-objective optimization discussed in Section 5.

were then used as the input to the DES. Based on the sample size analysis (see Section 3.5), ten runs of the DES were required to detect mean 90th percentile differences of at least 30 seconds below that of the current station. The mean and 95% confidence intervals were computed for the simulated 90th percentile response time across these ten runs, which are shown in Figure 4b. Each bar corresponds to a unique location proposed by the optimization procedure, and the size of each dot corresponds to the proportion of the 100 procedures that resulted in a location with the same 90th percentile response time based on the metamodel.

Based on this plot, the proposed station locations from the metamodel optimization built on all considered sample sizes reduced the simulated 90th percentile response time by at least 30 seconds compared to the current station. Note that for the sample size of 50, 1 optimization run resulted in a location with an average 90th percentiles response time of 142 minutes which is not included in Figure 4b. The best location, on average, was found by one optimization procedure with a sample size of 200, for a mean 90th percentile response time of 17.3 minutes compared to the current station's mean of 19.1 minutes. The next best location was found by 80% of the optimization procedures on the metamodel built on 1,000 locations, with a mean 90th percentile response time of 17.4 minutes.

In addition to the 90th percentile, Figure 5a shows the distribution of response times for several of these proposed locations as well as the current location from one run of the DES. While all distributions are right skewed, the current station's response times distribution has a peak near 3.8 minutes. The other curves, shown for sample sizes of 100, 400, and 1,000, have peaks between 4.5 and 6.5 minutes. This indicates that although the proposed locations reduce the simulated 90th percentile of response times, effectively reducing the amount of time needed to reach calls in outlying areas, they require additional time for many of the calls that are within five minutes of a response under the current system.

To understand the results described above, Figure 6a plots several of the proposed locations in relation to the city of Virginia. The current station, shown in red, is in the center of the city, allowing for a fast response to many calls in a dense area. However, the proposed locations based on the metamodel optimization are on the western side of the city, along US 53 or US 169, which allow for quicker access to the outlying regions. Thus, while there is some variability in the simulated response time distributions and 90th percentiles, the general trend suggests that moving the station along the highway on the western edge of the city will reduce the 90th percentile of the response times.

4.2. Two Station

Here results are presented for the case in which two stations are considered. Weighted random samples of two-station configurations were taken according to the sample sizes specified in Table 3, and these samples were used to fit and optimize station configurations using the metamodels. Additional details are provided in Section 9 of the Supplementary Material.

Figure 7a again shows the goodness of fit of the models by sample size using the cross-validated R^2 of the selected random forest metamodel along with standard error bars. Though the sample sizes are necessarily larger than the one-station version, an informative metamodel can still be achieved.

(a)

(b)

Figure 6: For the one-station setting, (a) location of proposed configurations in relation to Virginia; and (b) location of proposed configurations from the the multiobjective median-90th percentile optimization described in the Discussion. See Section 5.

The results of the optimization are not as clean here, with more diversity of locally optimal solutions for each sample size. Validating each of these metamodel-proposed solutions with ten runs on the DES found that, with the exception of a few outliers, nearly all proposed configurations improved upon the current system (Figure 7b). The best configuration was found by two optimization procedures on the metamodel built on a sample of 5,000 configurations, with a average 90th percentile of 13.5 minutes.

Figure 8 shows the distributions of simulated response times for the best configurations indicated by the metamodel with sample sizes of 1,000, 3,000, and 5,000. As with the one-station setting, while these proposed configurations decrease the 90th percentile, they increase the lower quantiles of response times. The peak of the distributions for each proposed configuration represented is shifted to the right of the current configuration's peak, signifying longer response times for many of the calls that are within three minutes of a response under the current system.

To understand this trend, Figure 9 shows the geographic locations of the most commonly proposed configuration for each sample size. The current configuration is shown in red, with locations in the center of Virginia and Eveleth. Each proposed configuration consisted of one station in the northern area to the west of Virginia and one station in the southern region to the southeast of Eveleth. All proposed configurations have locations closer to the highway than the current configuration, with the distances between stations slightly farther than the current distance. Similar to the one-station case, this general trend suggests that moving the stations closer to the highway would reduce the 90th percentile of the response times.

4.3. Twelve Stations

Here, briefer results are provided for the twelve-station case. Additional details for the twelve-station DES are in Supplementary Material Section 11, while results for an intermediate five-station case are provided in Section 10 of the Supplementary Material. Overall, the metamodels for this scenario needed much more data than the one- and two-station systems, while at the same time requiring much larger computational resources to simulate and optimize. The twelve-station case, and even the five-station case, were challenging scenarios that stretched this methodology and computational infrastructure, and consequently, the results are less encouraging. Still, it is important to demonstrate both where the metamodeling approach excels and where it requires more development.

For twelve stations, three different sizes of configuration samples were generated using the constrained, weighted method. Figure 10a shows the results; note that even with a sample size of 50,000, the cross-validated R^2 is not even 50%. PSO was then used to optimize the station locations using each metamodel, and the complexity of the problem resulted in a majority of the procedures converging to different configurations. Increasing the sample size used to train the metamodel did result in lower overall distributions of 90th percentiles for the 100 configurations, but in this case nearly all times were still greater than 30 minutes and none approached the current average of 18.5 minutes (Figure 10b). Clearly, the metamodels failed to adequately capture the complexities of the system. More data is needed to train the metamodels in this case.

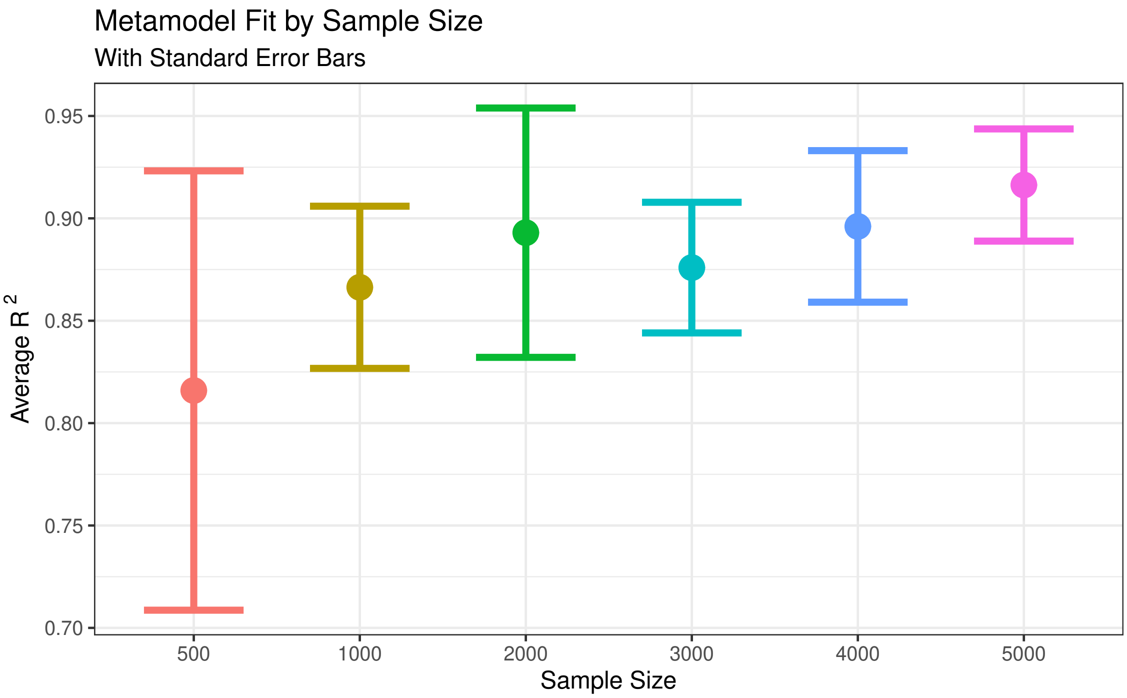

(a)

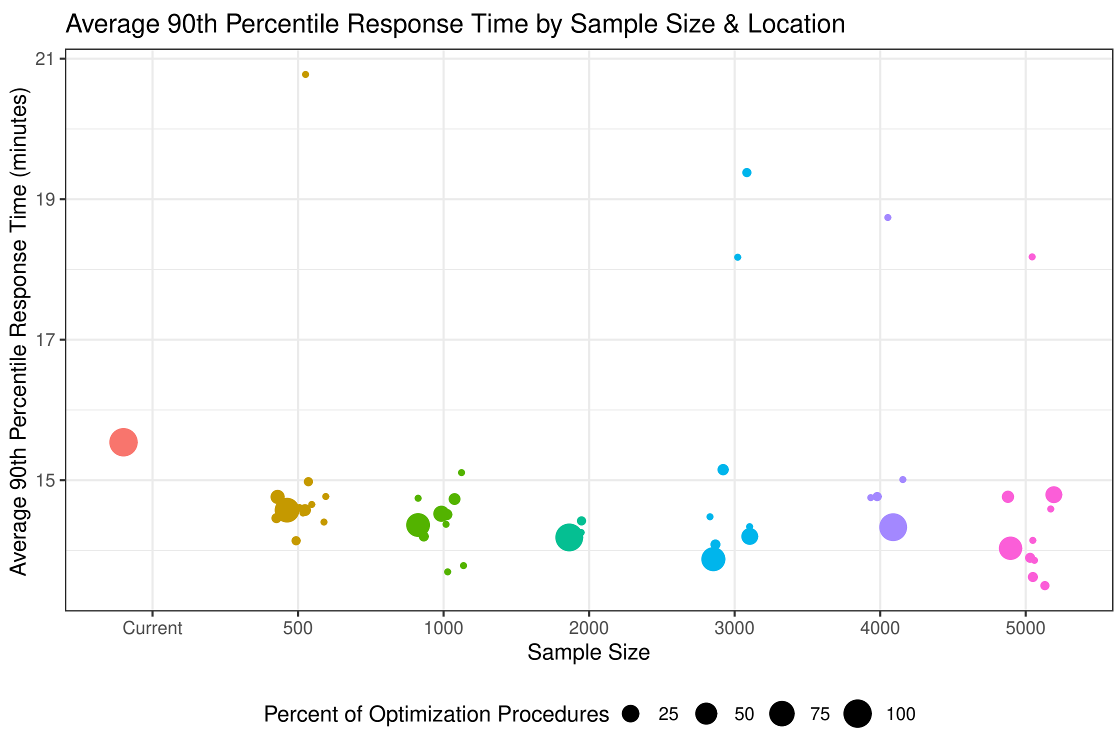

(b)

Figure 7: For the two-station setting, (a) metamodel fit by sample size with standard error bars; and (b) average DES simulated 90th percentile response times for proposed locations.

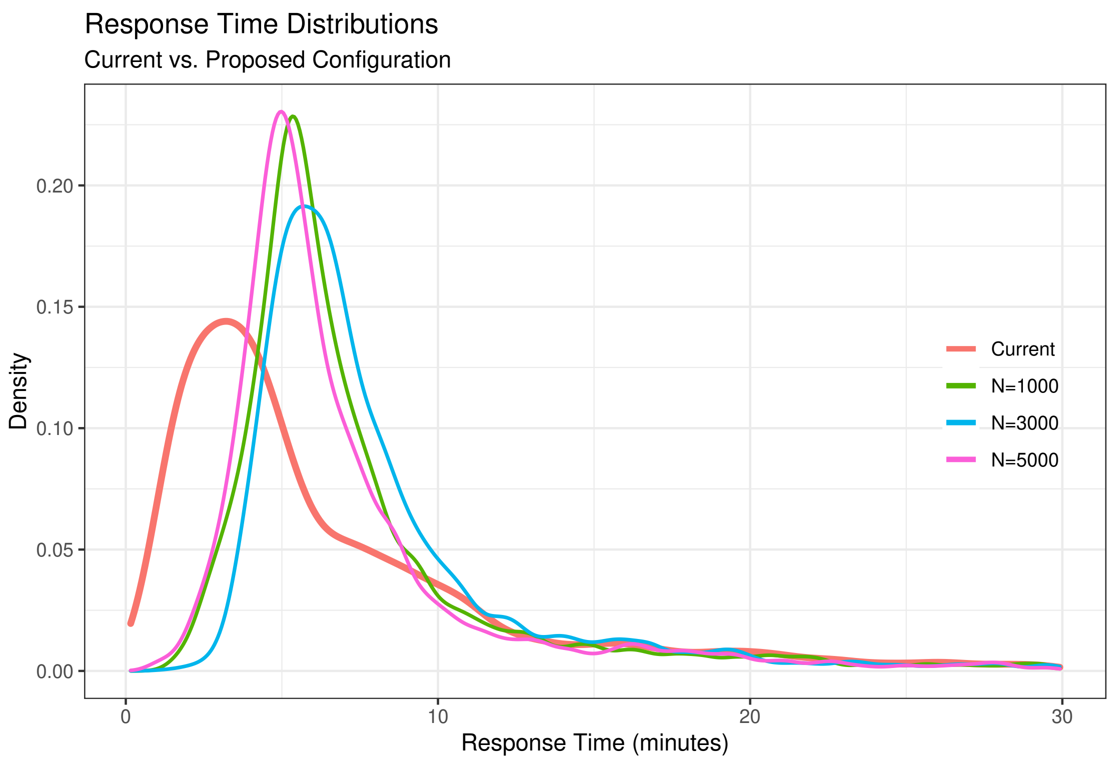

Figure 8: Distributions of response times for the current and proposed configurations, for the two-station setting.

5. Discussion and Conclusion

5.1. Discussion

Interestingly, the best, most reliable metamodel fits were not necessarily needed to find improved locations, and many models with poorer fits also proposed configurations with good average 90th percentile response times for the one- and two-station scenarios. This suggests an important point: metamodeling may not produce a single, globally optimal solution, but is effective at collecting a set of promising solutions that can then be more carefully evaluated and analyzed using a simulator and domain knowledge. This collection happens naturally when an optimization heuristic such as particle swarm optimization is used and produces a number of solutions based on optimization runs from different initializations. This set of solutions can then reveal what is in common among improved configurations. For instance, for the one- and two-station cases many of the solutions suggested by the optimization heuristic moved the stations closer to the highway (see Figures 6a and 9). This makes intuitive sense because the goal is to minimize the 90th percentile of the response time distribution; increases in the median response are allowed in order to reduce the right tail response by improving highway access (see Figures 5a and 8).

A simple question then presents itself: can the benefits of a reduced 90th percentile be achieved while also retaining a median response that is close to the current level? To briefly investigate this, a simple multiobjective optimization approach was used to simultaneously minimize the 90th percentile and median

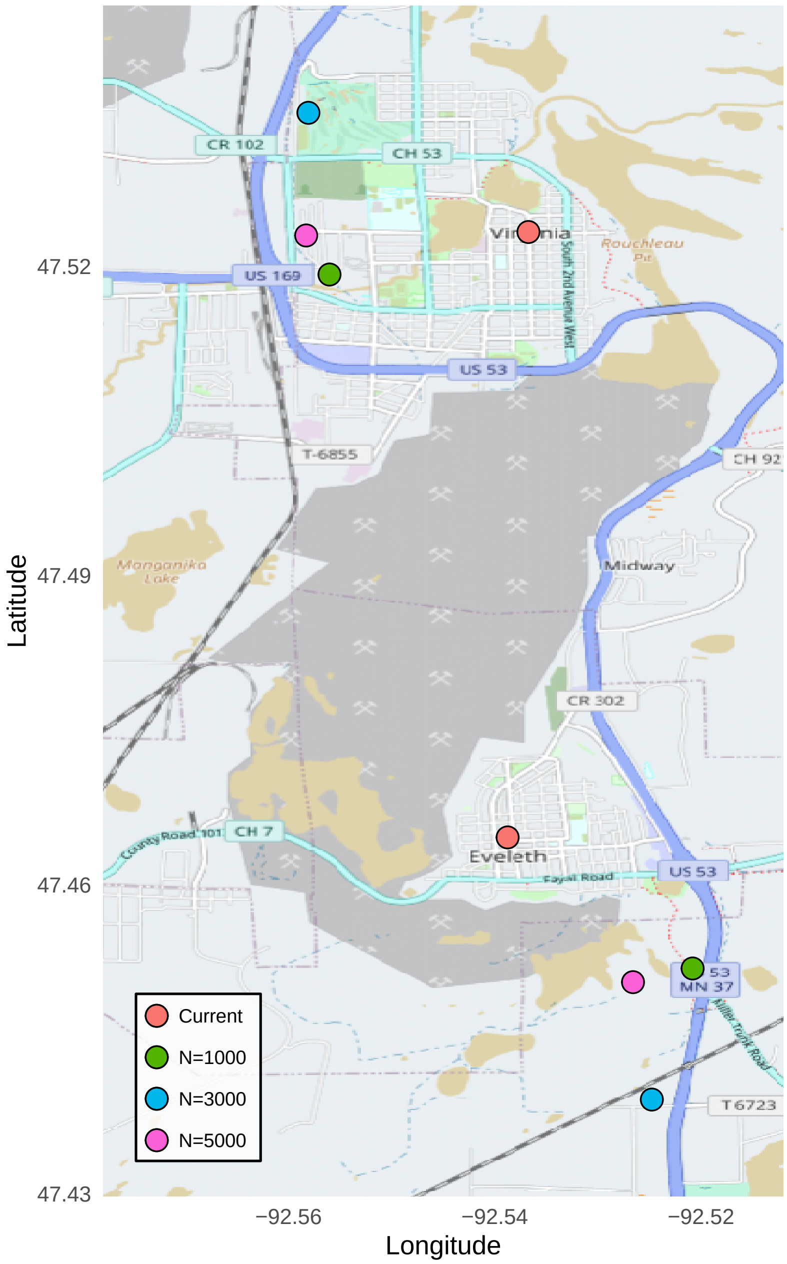

Figure 9: Location of proposed configurations in relation to St. Louis County, for the two-station setting.

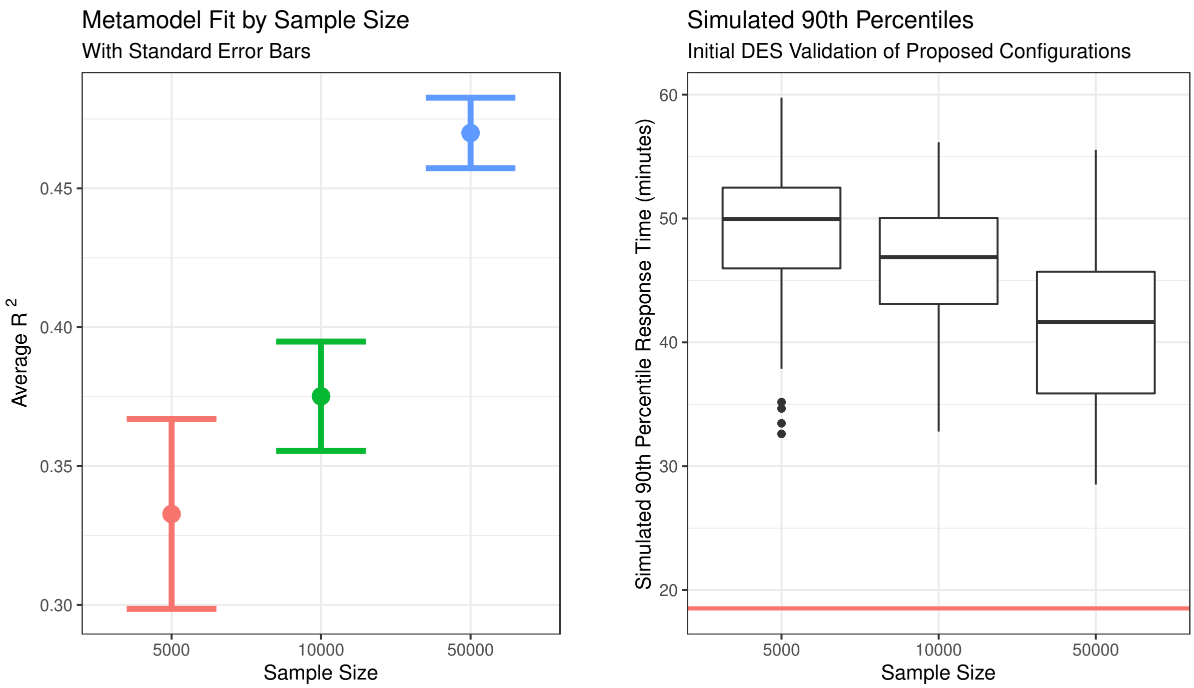

(a) (b)

Figure 10: For the twelve-station setting, (a) metamodel fit by sample size, with standard error bars; and (b) boxplot of single DES validation of 100 proposed configurations produced by optimizing the metamodel-predicted 90th percentile response time. The red line in (b) is the DES performance of the current configuration, with 95% confidence band.

response times for the one-station case. In this optimization, a new objective function was created, as a linear combination of 90th percentile and median, where both objectives were equally weighted. The analysis was done for sample sizes of 100, 200, 300, 400, and 500. Due to the increased stability around the median quantile, the R^2 for these metamodels were typically 5-10 percentages points higher than those modeling the 90th percentile. After the optimization and validation steps, most proposed locations decreased the 90th percentile by approximately 20 to 60 seconds without drastically affecting the median. Figure 5b shows the response time distributions for the current station location in red, proposed locations based on the single objective analysis in blue, and the proposed locations based on the multiobjective analysis in green using a sample size of 300; the multiple density curves within each analysis type represent the different locations suggested by the 100 different optimization procedures. Geographically, many of the proposed locations under the multiobjective optimization procedures were on the western edge of the city of Virginia, but not as close to the highways as seen in Figure 6a, which balances access to the outlying areas with proximity to the high density area. These proposed locations are seen in Figure 6b, color coded by the sample size used to build the underlying model.

We also explored the robustness of our work by performing two types of sensitivity analyses. First, for the two-station case, the stations in the configuration which was selected based on the optimization of the 5,000-point metamodel (shown in pink in Figure 9) were perturbed and reevaluated by both the metamodel and DES in order to assess geographical trends in the 90th percentile of the response times. Details are reported in Section 12.1 of the Supplementary Material, but overall the findings were as expected: perturbed locations close to the optimized locations had more similar predicted response times than those locations that were further away.

The second sensitivity analysis assessed the robustness of the methods to small changes in the gamma and lognormal distributional parameters used in the DES (see Section 3.2). For the work in this paper, the parameters in these distributions were estimated from the data and as such include measured uncertainty. For this sensitivity analysis, random perturbations of these parameters were generated based on their estimates and their estimated uncertainty, and these new parameters were used to create a new version of the DES. This process was repeated 50 times to create 50 different simulators. Each perturbed DES was then used to produce a set of training data used to fit a metamodel and optimize the station location using the methods described in this paper. Additional details are provided in Section 12.2 of the Supplementary Material. Figure 11 displays a map with the current station location in red, the original proposed location in blue, with the 50 proposed locations based on the perturbed simulators in green. While the specific locations vary, the message is clear and consistent with earlier findings: to minimize the 90th percentile of the response times in the one-station setting, the ideal station location is along the highway near the intersection of US 53 and US 169. The results of this sensitivity analysis suggest that slightly different input parameters to the DES, as could be expected from utilizing a different dataset from the same system, still yield consistent qualitative suggestions for the EMS system.

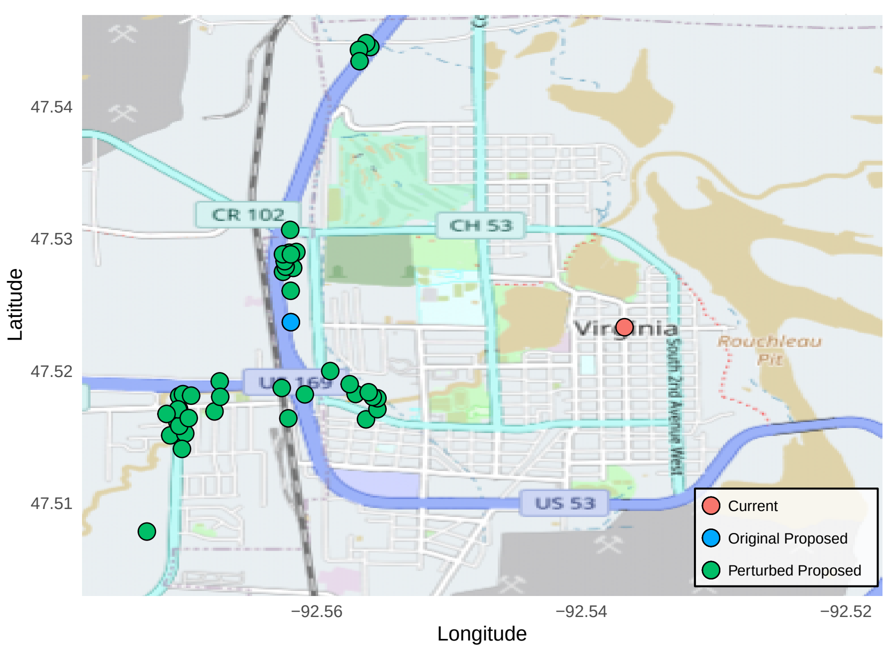

Figure 11: Location of the proposed station locations using the simulators with perturbed parameters in relation to the current station and original proposed station location.

5.2. Conclusion

Because detailed discrete event simulators for large rural-urban EMS systems are computationally demanding, it is difficult to optimize using the DES directly, and even more challenging when including randomness in the simulation. To address this problem, metamodels are used as a surrogate for the DES and the present work examined how much data these surrogates required as the complexity of the system increased. The study found that in order to model the 90th percentile response time of a single station, only 200 to 400 data points are needed to obtain an accurate random forest metamodel. This metamodel, when used to optimized, produced configurations that improved upon the current configuration by over a minute, according to the DES. For the two-station setting, around 3,000 data points were required, and resulting configurations suggested improvements of at least 30 seconds. For the twelve-station setting, however, even 50,000 data points was inadequate to train a metamodel which reliably optimized the system. Thus, while using a constrained, weighted sampling technique to generate data in intuitively promising regions decreased the number of DES runs needed, the metamodeling approach still requires a significant amount of time and data for complicated settings. These findings, it is expected, can be generalized to similar rural-uban EMS systems, with the quality of the metamodel—as measured, say, by out-of-sample R^2—serving as a rough indicator of how effective the optimization will be.

Based on this analysis, the discrete event simulators constructed to model EMS systems in St. Louis County, Minnesota appear to be reasonable approximations of the system's responses to medical calls. The use of randomness in each aspect of the simulator, guided by historic call data, strengthens the reliability of the DES by ensuring that anomalous calls do not drastically alter a given station configuration's responses. However, these simulators are overly simplified representations of the true system and do not account for the many complexities. For instance, all calls are assumed to have the same priority and only medical calls were considered, even though resources at these stations are shared between medical and fire emergencies. In addition, it is assumed that only one vehicle is needed for each call, and this responding vehicle is an ambulance that always starts from the station, not the hospital, a previous call location, or any location in between. Thus, while this DES can be used to find improved configurations, a more complex simulator could provide more insight.

There are many areas of future work possible to continue or improve this analysis. To start, developing a faster simulator would allow for a fuller understanding of the amount of data required for the more complex cases with many stations. Additionally, rather than building a single metamodel on many runs of the DES, exploring model averaging techniques for many metamodels built on smaller samples could provide valuable insight with fewer runs. Alternatively, since an EMS system is unlikely to have the capacity to move all stations in practice, an alternative approach that focuses on either moving a subset of the stations or removing certain stations, while keeping the others in place, could be considered. Another approach to reduce the amount of data necessary to fit quality metamodels would be to continue exploring alternative sampling strategies. Techniques incorporating spatial densities and inverse-distance weights could be further explored, in addition to adaptive sampling, also known as active learning, which has been shown in the metamodeling literature to produce small yet informative samples [39, 41]. Also, several aspects regarding the optimization could be investigated. Though PSO appears to be a solid metaheuristic choice, it is possible that other methods, such as Genetic Algorithms or tabu search, would improve results for a given level of metamodel quality. Along with the continuous optimization of station locations, incorporating station resources such as vehicles or personnel in the optimization would allow for more thorough solutions; for instance, perhaps combining the resources of two stations into a single station could reduce the response times while also reducing cost. Finally, performing a full multiobjective optimization analysis with Pareto fronts is a promising area of future research, as discussed above. Invariably, EMS departments care about more than one measure, and a multiobjective approach will provide decision-makers with a more realistic menu of options to improve their systems.

6. Acknowledgements

The authors are grateful to Chief Allen Lewis and the Virginia, Minnesota Fire Department for providing the data for this analysis and for their guidance with the standards and structures of EMS systems. They also thank Levrum Data Technologies for their advice and direction regarding the motivation for this project. In addition, Jens Mueller, Director of High Performance Computing Services, and Noah Dunn, Master's student in Computer Science, both of Miami University, provided invaluable computing and networking assistance.

7. Funding Sources and Competing Interests

This research did not receive any specific grant funding from agencies in the public, commercial, or not-for-profit sectors. Declarations of interest: none.

References

[1] L. Aboueljinane, E. Sahin, Z. Jemai, A review on simulation models applied to emergency medical service operations, Computers & Industrial Engineering (2013) 734–750doi:https://doi.org/10.1016/j.cie.2013.09.017.
URL `https://www.sciencedirect.com/science/article/pii/S0360835213003100`

[2] Z. He, X. Qin, Y. Xie, J. Guo, Service Location Optimization Model for Improving Rural Emergency Medical Services, Transportation Research Record: Journal of the Transportation Research Board 2672 (32) (2018) 83–93. doi:10.1177/0361198118791363.
URL `http://journals.sagepub.com/doi/10.1177/0361198118791363`

[3] NFPA 1710: Standard for the Organization and Deployment of Fire Suppression Operations, Emergency Medical Operations, and Special Operations to the Public by Career Fire Departments, NFPA, Quincy, MA, 2020.

[4] H. Andersson, T. A. Granberg, M. Christiansen, E. S. Aartun, H. Leknes, Using optimization to provide decision support for strategic emergency medical service planning – Three case studies, International Journal of Medical Informatics (2020). doi:10.1016/j.ijmedinf.2019.103975.

[5] M. I. H. Umam, B. Santosa, N. Siswanto, Minimizing response time in Medical Emergency Service: A literature review, Proceedings of the International Conference on Industrial Engineering and Operations Management 2018-March (2014) (2018) 2511–2520.

[6] M. A. Zaffar, H. K. Rajagopalan, C. Saydam, M. Mayorga, E. Sharer, Coverage, survivability or response time: A comparative study of performance statistics used in ambulance location models via simulation–optimization, Operations Research for Health Care (2016). doi:10.1016/j.orhc.2016.08.001.

[7] J. Yao, X. Zhang, A. T. Murray, Location optimization of urban fire stations: Access and service coverage, Computers, Environment and Urban Systems (2019). doi:10.1016/j.compenvurbsys.2018.10.006.

[8] M. Liu, D. Yang, F. Hao, Optimization for the Locations of Ambulances under Two-Stage Life Rescue in the Emergency Medical Service: A Case Study in Shanghai, China, Mathematical Problems in Engineering (2017). doi:10.1155/2017/1830480.

[9] E. S. Savas, Simulation and Cost-Effectiveness Analysis of New York's Emergency Ambulance Service, Management Science (1969). doi:10.1287/mnsc.15.12.b608.

[10] X. Li, Z. Zhao, X. Zhu, T. Wyatt, Covering models and optimization techniques for emergency response facility location and planning: a review, Mathematical Methods of Operations Research 74 (3) (2011) 281–310. doi:10.1007/s00186-011-0363-4.
URL http://link.springer.com/10.1007/s00186-011-0363-4

[11] R. Aringhieri, M. Bruni, S. Khodaparasti, J. van Essen, Emergency medical services and beyond: Addressing new challenges through a wide literature review, Computers & Operations Research 78 (2017) 349–368. doi:10.1016/j.cor.2016.09.016.
URL https://linkinghub.elsevier.com/retrieve/pii/S0305054816302362

[12] K. Degeling, M. IJzerman, H. Koffijberg, A scoping review of metamodeling applications and opportunities for advanced health economic analyses, Expert Review of Pharmacoeconomics & Outcomes Research (2019) 181–187PMID: 30426801. arXiv:https://doi.org/10.1080/14737167.2019.1548279, doi:10.1080/14737167.2019.1548279.
URL https://doi.org/10.1080/14737167.2019.1548279

[13] P. Lutter, D. Degel, L. Wiesche, B. Werners, Analysis of Ambulance Location Models Using Discrete Event Simulation, in: M. Lübbecke, A. Koster, P. Letmathe, R. Madlener, B. Peis, G. Walther (Eds.),

Operations Research Proceedings, Springer, Cham, 2016, pp. 377–383. doi:10.1007/978-3-319-28697-6_53.
URL http://link.springer.com/10.1007/978-3-319-28697-6_53

[14] M. van Buuren, G. J. Kommer, R. van der Mei, S. Bhulai, EMS call center models with and without function differentiation: A comparison, Operations Research for Health Care 12 (2017) 16–28. doi:10.1016/j.orhc.2016.12.001.

[15] T. Ünlüyurt, Y. Tunçer, Estimating the performance of emergency medical service location models via discrete event simulation, Computers and Industrial Engineering (2016). doi:10.1016/j.cie.2016.03.029.

[16] M. Allen, A. Spencer, A. Gibson, J. Matthews, A. Allwood, S. Prosser, M. Pitt, Right cot, right place, right time: improving the design and organisation of neonatal care networks – a computer simulation study, Health Services and Delivery Research (2015). doi:10.3310/hsdr03200.

[17] M. A. Ahmed, T. M. Alkhamis, Simulation optimization for an emergency department healthcare unit in Kuwait, European Journal of Operational Research (2009). doi:10.1016/j.ejor.2008.10.025.

[18] S. Wu, Agent-based discrete event simulation modeling and evolutionary real-time decision making for large-scale systems, University of Pittsburgh (2008).

[19] J. Karnon, J. Stahl, A. Brennan, J. J. Caro, J. Mar, J. Möller, Modeling using discrete event simulation: A report of the ISPOR-SMDM modeling good research practices task force-4, Medical Decision Making (2012). doi:10.1177/0272989X12455462.

[20] J. S. Barrett, B. Jayaraman, D. Patel, J. M. Skolnik, A SAS-based solution to evaluate study design efficiency of phase I pediatric oncology trials via discrete event simulation, Computer Methods and Programs in Biomedicine (2008). doi:10.1016/j.cmpb.2007.12.008.

[21] Z. Abdmouleh, A. Gastli, L. Ben-Brahim, M. Haouari, N. A. Al-Emadi, Review of optimization techniques applied for the integration of distributed generation from renewable energy sources, Renewable Energy (2017) 266–280doi:https://doi.org/10.1016/j.renene.2017.05.087.
URL https://www.sciencedirect.com/science/article/pii/S0960148117304822

[22] S. Ridler, A. J. Mason, A. Raith, A simulation and optimisation package for emergency medical services, European Journal of Operational Research (xxxx) (2021). doi:10.1016/j.ejor.2021.07.038.
URL https://doi.org/10.1016/j.ejor.2021.07.038

[23] R. Aringhieri, An integrated DE and AB simulation model for EMS management, 2010 IEEE Workshop on Health Care Management, WHCM 2010 (2010). doi:10.1109/WHCM.2010.5441260.

[24] Y. Wang, K. L. Luangkesorn, L. Shuman, Modeling emergency medical response to a mass casualty incident using agent based simulation, Socio-Economic Planning Sciences 46 (4) (2012) 281–290.

doi:10.1016/j.seps.2012.07.002.
URL http://dx.doi.org/10.1016/j.seps.2012.07.002

[25] A. Anagnostou, A. Nouman, S. J. Taylor, Distributed hybrid agent-based discrete event emergency medical services simulation, Proceedings of the 2013 Winter Simulation Conference - Simulation: Making Decisions in a Complex World, WSC 2013 (2013) 1625–1636doi:10.1109/WSC.2013.6721545.

[26] A. Ingolfsson, E. Erkut, S. Budge, Simulation of single start station for Edmonton EMS, Journal of the Operational Research Society (2003). doi:10.1057/palgrave.jors.2601574.

[27] S. S. Wei Lam, Z. C. Zhang, H. C. Oh, Y. Y. Ng, W. Wah, M. E. Hock Ong, Reducing ambulance response times using discrete event simulation, Prehospital Emergency Care (2014). doi:10.3109/10903127.2013.836266.

[28] L. Aboueljinane, Z. Jemai, E. Sahin, Reducing ambulance response time using simulation: The case of Val-de-Marne department Emergency Medical service, in: Proceedings - Winter Simulation Conference, 2012. doi:10.1109/WSC.2012.6465018.

[29] N.-H. Thi Nguyen, Quantitative Analysis of Ambulance Location-allocation and Ambulance State Prediction, Ph.D. thesis, Linköping University Electronic Press (2015). doi:10.3384/lic.diva-113346.

[30] J. Schneider, M. Schröder, Simulation-Based Location Optimization of Ambulance Stations, in: B. Fortz, M. Labbé (Eds.), Operations Research Proceedings 2018, Springer, Cham, 2019, pp. 143–149. doi:10.1007/978-3-030-18500-8_19.
URL http://link.springer.com/10.1007/978-3-030-18500-8_19

[31] A. J. Mason, Simulation and Real-Time Optimised Relocation for Improving Ambulance Operations, 2013, pp. 289–317. doi:10.1007/978-1-4614-5885-2_11.
URL http://link.springer.com/10.1007/978-1-4614-5885-2_11

[32] H. Leknes, E. S. Aartun, H. Andersson, M. Christiansen, T. A. Granberg, Strategic ambulance location for heterogeneous regions, European Journal of Operational Research 260 (1) (2017) 122–133. doi:10.1016/j.ejor.2016.12.020.

[33] V. A. Knight, P. R. Harper, L. Smith, Ambulance allocation for maximal survival with heterogeneous outcome measures, Omega 40 (6) (2012) 918–926. doi:10.1016/j.omega.2012.02.003.
URL http://dx.doi.org/10.1016/j.omega.2012.02.003

[34] E. Erkut, A. Ingolfsson, G. Erdogan, Ambulance location for maximum survival, Naval Research Logistics 55 (1) (2008) 42–58. doi:10.1002/nav.20267.

[35] M. van Buuren, R. van der Mei, S. Bhulai, Demand-point constrained EMS vehicle allocation problems for regions with both urban and rural areas, Operations Research for Health Care 18 (2018) 65–83.

doi:10.1016/j.orhc.2017.03.001.
URL http://dx.doi.org/10.1016/j.orhc.2017.03.001

[36] S. Chanta, M. E. Mayorga, L. A. McLay, Improving emergency service in rural areas: a bi-objective covering location model for EMS systems, Annals of Operations Research 221 (1) (2014) 133–159. doi:10.1007/s10479-011-0972-6.

[37] S. Felder, H. Brinkmann, Spatial allocation of emergency medical services: Minimising the death rate or providing equal access?, Regional Science and Urban Economics 32 (1) (2002) 27–45. doi:10.1016/S0166-0462(01)00074-6.

[38] M. Amorim, F. Antunes, S. Ferreira, A. Couto, An integrated approach for strategic and tactical decisions for the emergency medical service: Exploring optimization and metamodel-based simulation for vehicle location, Computers and Industrial Engineering (2019). doi:10.1016/j.cie.2019.106057.

[39] F. Antunes, M. Amorim, F. C. Pereira, B. Ribeiro, Active learning metamodeling for policy analysis: Application to an emergency medical service simulator, Simulation Modelling Practice and Theory (2019). doi:10.1016/j.simpat.2019.101947.

[40] S. Dey, T. Mukhopadhyay, S. Adhikari, Metamodel based high-fidelity stochastic analysis of composite laminates: A concise review with critical comparative assessment, Composite Structures (2017) 227–250doi:https://doi.org/10.1016/j.compstruct.2017.01.061.
URL https://www.sciencedirect.com/science/article/pii/S0263822316328793

[41] H. Liu, Y. Ong, J. Cai, A survey of adaptive sampling for global metamodeling in support of simulation-based complex engineering design, Structural and Multidisciplinary Optimization (2018) 393–416doi:https://doi.org/10.1007/s00158-017-1739-8.

[42] B. Can, C. Heavey, A comparison of genetic programming and artificial neural networks in metamodeling of discrete-event simulation models, Computers and Operations Research (2012). doi:10.1016/j.cor.2011.05.004.

[43] Y. F. Li, S. H. Ng, M. Xie, T. N. Goh, A systematic comparison of metamodeling techniques for simulation optimization in Decision Support Systems, Applied Soft Computing Journal (2010). doi:10.1016/j.asoc.2009.11.034.

[44] N. Villa-Vialaneix, M. Follador, M. Ratto, A. Leip, A comparison of eight metamodeling techniques for the simulation of N 2O fluxes and N leaching from corn crops, Environmental Modelling and Software (2012). doi:10.1016/j.envsoft.2011.05.003.

[45] F. A. Viana, T. W. Simpson, V. Balabanov, V. Toropov, Metamodeling in multidisciplinary design optimization: How far have we really come?, in: AIAA Journal, 2014. doi:10.2514/1.J052375.

[46] R. R. Barton, M. Meckesheimer, Chapter 18 Metamodel-Based Simulation Optimization (2006). doi:10.1016/S0927-0507(06)13018-2.

[47] R. R. Barton, Simulation optimization using metamodels, in: Proceedings - Winter Simulation Conference, 2009. doi:10.1109/WSC.2009.5429328.

[48] R. Hopkins, B. Smucker, Metamodeling to Improve Emergency Medical Systems: Exploring Two Case Studies, unpublished (2020).

[49] R. Jin, W. Chen, T. Simpson, Comparative studies of metamodelling techniques under multiple modelling criteria, Structural and Multidisciplinary Optimization 23 (1) (2001) 1–13. doi:10.1007/s00158-001-0160-4.
URL http://link.springer.com/10.1007/s00158-001-0160-4

[50] R. J. Yang, N. Wang, C. H. Tho, J. P. Bobineau, B. P. Wang, Metamodeling Development for Vehicle Frontal Impact Simulation, Journal of Mechanical Design 127 (5) (2005) 1014. doi:10.1115/1.1906264.
URL http://mechanicaldesign.asmedigitalcollection.asme.org/article.aspx?articleid=1448673

[51] B.-S. Kim, Y.-B. Lee, D.-H. Choi, Comparison study on the accuracy of metamodeling technique for non-convex functions, Journal of Mechanical Science and Technology 23 (4) (2009) 1175–1181. doi:10.1007/s12206-008-1201-3.
URL http://link.springer.com/10.1007/s12206-008-1201-3

[52] M. R. Kianifar, F. Campean, Performance evaluation of metamodelling methods for engineering problems: towards a practitioner guide, Structural and Multidisciplinary Optimization 61 (1) (2020) 159–186. doi:10.1007/s00158-019-02352-1.
URL http://link.springer.com/10.1007/s00158-019-02352-1

[53] T. Østergård, R. L. Jensen, S. E. Maagaard, A comparison of six metamodeling techniques applied to building performance simulations, Applied Energy (2018). doi:10.1016/j.apenergy.2017.10.102.

[54] C. Osorio, L. Chong, A Computationally Efficient Simulation-Based Optimization Algorithm for Large-Scale Urban Transportation Problems, Transportation Science 49 (3) (2015) 623–636. doi:10.1287/trsc.2014.0550.
URL http://pubsonline.informs.org/doi/10.1287/trsc.2014.0550

[55] J. Zentner, V. Kumar, D. Mavris, A Hierarchical Metamodeling Method for Large Scale, Multi-Objective Computer Simulations, in: 50th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, American Institute of Aeronautics and Astronautics, Reston, Virigina, 2009. doi:10.2514/6.2009-2238.
URL https://arc.aiaa.org/doi/10.2514/6.2009-2238

[56] Y. Ju, C. Zhang, L. Ma, Artificial intelligence metamodel comparison and application to wind turbine airfoil uncertainty analysis, Advances in Mechanical Engineering (2016). doi:10.1177/1687814016647317.

[57] L. Van Gelder, P. Das, H. Janssen, S. Roels, Comparative study of metamodelling techniques in building energy simulation: Guidelines for practitioners, Simulation Modelling Practice and Theory 49 (2014) 245–257. doi:10.1016/j.simpat.2014.10.004.
URL https://linkinghub.elsevier.com/retrieve/pii/S1569190X14001555

[58] F. M. Alam, K. R. McNaught, T. J. Ringrose, A comparison of experimental designs in the development of a neural network simulation metamodel, Simulation Modelling Practice and Theory 12 (7-8) (2004) 559–578. doi:10.1016/j.simpat.2003.10.006.
URL https://linkinghub.elsevier.com/retrieve/pii/S1569190X04000747

[59] L. Ding, X. Zhang, Sample and Computationally Efficient Simulation Metamodeling in High Dimensions (oct 2020). arXiv:2010.06802.
URL http://arxiv.org/abs/2010.06802

[60] F. Zeinali, M. Mahootchi, M. M. Sepehri, Resource planning in the emergency departments: A simulation-based metamodeling approach, Simulation Modelling Practice and Theory (2015). doi:10.1016/j.simpat.2015.02.002.

[61] R Core Team, R: A Language and Environment for Statistical Computing (2020).
URL https://www.r-project.org/

[62] H. Wickham, M. Averick, J. Bryan, W. Chang, L. McGowan, R. François, G. Grolemund, A. Hayes, L. Henry, J. Hester, M. Kuhn, T. Pedersen, E. Miller, S. Bache, K. Müller, J. Ooms, D. Robinson, D. Seidel, V. Spinu, K. Takahashi, D. Vaughan, C. Wilke, K. Woo, H. Yutani, Welcome to the Tidyverse, Journal of Open Source Software 4 (43) (2019) 1686. doi:10.21105/joss.01686.
URL https://joss.theoj.org/papers/10.21105/joss.01686

[63] I. Ucar, B. Smeets, A. Azcorra, simmer : Discrete-Event Simulation for R, Journal of Statistical Software 90 (2) (2019) 1–30. doi:10.18637/jss.v090.i02.
URL http://www.jstatsoft.org/v90/i02/

[64] T. Giraud, osrm: Interface Between R and the OpenStreetMap-Based Routing Service OSRM (2020).
URL https://cran.r-project.org/package=osrm

[65] M. Kuhn, caret: Classification and Regression Training (2020).
URL https://cran.r-project.org/package=caret

[66] A. Liaw, M. Wiener, Classification and Regression by randomForest, R News 2 (3) (2002) 18–22.
URL https://cran.r-project.org/doc/Rnews/

[67] C. Bendtsen, pso: Particle Swarm Optimization (2012).
URL https://cran.r-project.org/package=pso

[68] A. Baddeley, E. Rubak, R. Turner, Spatial Point Patterns: Methodology and Applications with R, Chapman and Hall/CRC Press, 2015.
URL http://www.crcpress.com/Spatial-Point-Patterns-Methodology-and-Applications-with-R/Baddeley-Rubak-Turner/9781482210200/

[69] R. McCormack, G. Coates, A simulation model to enable the optimization of ambulance fleet allocation and base station location for increased patient survival, European Journal of Operational Research (2015). doi:10.1016/j.ejor.2015.05.040.

[70] M. Hatta, C. S. Lim, A. Faiz, Z. Abidin, M. H. Azizan, S. S. Teoh, Solving maximal covering location with particle swarm optimization, International Journal of Engineering and Technology 5 (2013) 3301–3306.

[71] S. Twaha, M. A. Ramli, A review of optimization approaches for hybrid distributed energy generation systems: Off-grid and grid-connected systems, Sustainable Cities and Society (2018) 320–331doi:https://doi.org/10.1016/j.scs.2018.05.027.
URL https://www.sciencedirect.com/science/article/pii/S2210670718301264

[72] S. K. ElKady, H. M. Abdelsalam, A Modified Particle Swarm Optimization Algorithm for Solving Capacitated Maximal Covering Location Problem in Healthcare Systems, 2016, pp. 117–133. doi:10.1007/978-3-319-21212-8_5.
URL http://link.springer.com/10.1007/978-3-319-21212-8_5

[73] S. Panda, N. P. Padhy, Comparison of particle swarm optimization and genetic algorithm for FACTS-based controller design, Applied Soft Computing Journal 8 (4) (2008) 1418–1427. doi:10.1016/j.asoc.2007.10.009.