Forecasting the daily demand for emergency medical ambulances in England and Wales: A benchmark model and external validation

Thomas Monks¹, Michael Allen², Alison Harper², Andrew Mayne³, and Lucy Collins³

¹University of Exeter Medical School

²University of Exeter Medical School & NIHR South West Peninsula Applied Research Collaboration (ARC).

³South Western Ambulance Foundation Trust

May 28, 2021

Abstract

Background We aimed to select and externally validate a benchmark method for emergency ambulance services to use to forecast the daily number of calls that result in the dispatch of one or more ambulances. The study was conducted using standard methods known to the UK's NHS to aid implementation in practice.

Methods We selected our benchmark model from a naive benchmark and 14 standard forecasting methods. Mean absolute scaled error and 80 and 95% prediction interval coverage over a 84 day horizon were evaluated using time series cross validation across eight time series from the South West of England. External validation was conducted by time series cross validation across 13 time series from London, Yorkshire and Welsh Ambulance Services.

Results A model combining a simple average of Facebook's Prophet and regression with ARIMA Errors (1, 1, 3)(1, 0, 1, 7) was selected. Benchmark MASE, 80 and 95% prediction intervals were 0.68 (95% CI 0.67 - 0.69), 0.847 (95% CI 0.843 - 0.851), and 0.965 (95% CI 0.949 - 0.977), respectively. Performance in the validation set was within expected ranges for MASE, 0.73 (95% CI 0.72 - 0.74) 80% coverage (0.833; 95% CI 0.828-0.838), and 95% coverage (0.965; 95% CI 0.963-0.967).

Conclusions We provide a robust externally validated benchmark for future ambulance demand forecasting studies to improve on. Our benchmark forecasting model is high quality and usable by ambulance services. We provide a simple python framework to aid its implementation in practice.

1 Introduction

Ambulance response times can be critical to patient outcomes for serious clinical events such as cardiac arrest [1], stroke [2], and major trauma [3]. Managing ambulance provision efficiently is therefore critically important to health outcomes. Part of that management is having an accurate forecast of expected demand that can be used to plan and schedule appropriate workforce at the regional and daily level. Forecasting studies of demand for emergency medical services (EMS) date back over three decades [4, 5, 6, 7], yet there is minimal evidence of incremental improvement in methods, accuracy or crucial implementation in ambulance services. We argue that this stems from a number of limitations of the current evidence. A fundamental weakness is that studies are single site with no external validation of the forecasting methods chosen. It is not possible to directly compare forecast accuracy across settings as individual studies have opted for scale dependent forecast errors. Existing studies therefore provide no way for an EMS or researcher to judge if their forecasting methods are on par, fall below or exceed a state-of-the-art benchmark. There are two additional subtle limitations to the evidence. First, existing studies have tended to focus evaluation on point forecast accuracy; for example, the accuracy of a prediction of demand next Tuesday. In a statistical perspective on forecasting, a point forecast is accompanied by a prediction interval: a range of values that a future value might take with a high probability [8]. For example, next Tuesday’s prediction could be supplied with a 95% prediction interval stating that the actual value for next Tuesday should lie within a given range with probability 0.95. Achieving adequate prediction interval coverage is difficult and an understanding of a chosen method’s capability is of high importance. Second, the studies that have investigated novel high quality methods for forecasting EMS demand have often neglected to prove their method is more accurate than the most naive of benchmarks. For instance, using the last observed value as the forecast. A strong evidence base would make this comparison without assuming *a priori* that *any* method is more accurate.

This study aims to provide EMS forecasting benchmarks for future research to incrementally improve on. To aid EMS workforce planning, we predict the daily number of calls that result in the dispatch of one or more ambulances [5]. Our objectives are to establish statistical benchmarks for point forecast accuracy and prediction interval coverage up to 12 weeks (84 days) ahead. We select the most accurate forecasting method from 14 established methods and present details of the forecast error distribution at 7 day intervals. We then externally test the selected method in three further ambulance trusts (in the UK an ambulance trust provides services across a region). Overall we test the method in 21 time series. Finally, we also provide an easy to use MIT licensed (free and open) Python framework for deploying the model. Our results enable future replication studies or studies of new methods to directly compare their accuracy to our benchmark. EMS data science teams developing in-house forecasting tools or purchasing commercial forecasting systems can also compare results and spend NHS resource wisely.

2 Methods

We build multiple time series forecasting models using established methods. Model selection is conducted by time series cross validation evaluating both point forecast accuracy and prediction interval coverage. All models are compared to a naive statistical baseline model. We then test the selected method in a simulated forecasting setting (seven regions within a NHS Trust). The accuracy of the method in our test set is our benchmark. We test the external validation of the benchmark by applying the method to a further 13 regions from three different ambulance trusts in the United Kingdom.

2.1 Study setting

We develop our benchmark model using data from the South West Ambulance Service Trust (SWAST). SWAST is an NHS Trust in England that provides emergency medical services to a population of 5.6 million spread over a mixed urban/rural region of 26,000 km². The service receives an average 2,300 calls per day that require the dispatch of one or more ambulances. Forecasts are rerun every week and are used to set the staffing rotas three months ahead. The South West region can be broken down into seven sub-regions / time series: Devon, Cornwall, Dorset, Somerset, Gloucestershire, Wiltshire, and ‘Bristol, North Somerset and South Gloucestershire (BNSSG)’.

Our external validation is comprised of data from London Ambulance Service (LAS; 5 time series; population 8.6m; area size 1,570km²), the Welsh Ambulance Service Trust (WAST; 4 time series; pop 3m,

area 20,740km²) and the Yorkshire Ambulance Service (YAS; 6 time series; pop 5m; area 15,540km²). Overall the trusts serve a population of over 20 million people in the UK.

2.2 Outcome measures

We measure both point forecast error and prediction interval coverage of the forecasts. For point forecast error, our main outcome measure is the *Mean Absolute Scaled Error (MASE)* as this provides an easy to understand relative error measure that can be compared across ambulance trusts. For MASE, we scale the out of sample *mean absolute error* by the equivalent one-step within-sample error from a Seasonal Naive model [9]. We also report two secondary point forecast error measures. A second relative error measure is the *symmetric Mean Absolute Percentage Error (sMAPE)*. As it has been used elsewhere [5], we also provide a scale-dependent measure via the *Root Mean Absolute Squared Error (RMSE)*, although this can only be used in the context of the specific time series. Prediction interval coverage measures the proportion of out-of-sample observations that fall within a prediction interval with an expected probability. For example, it is expected that 80% of points will fall within an 80% prediction interval. For cross-validation, we report the 80% and 95% prediction interval coverage. We report the full forecast distribution for the final benchmark model.

2.3 Data sources

The study was conducted at a daily time-series level. Each observation represents the total number of emergency calls that resulted in the dispatch of one or more emergency ambulances. For model development, each sub-region in the data were broken into training (01/01/13 - 30/06/17; n = 1279), validation (01/07/17 - 31/01/18; n = 549) and test sets (01/01/2019 - 31/12/2019; n = 365). The validation period differs from test in that it is used to tune and select models.

2.4 Analysis environment

All analysis code was written in Python 3.7.5 and R 3.6.1. Python forecasting libraries used were pmdarima v1.5.3 [10], fbprophet v0.5 [11], statsmodels v0.11.1 [12], tbats v1.0.10 [13], forecast-tools v0.1.5 [14] and from R we used Rssa v1.0.2 [15]. Data cleaning and manipulation were done using Pandas [16] and NumPy [17]. All charts were produced with Matplotlib [18]. To reproduce the results of the study, we provide all code and data, a Docker image containing the exact software environment, a conda virtual environment (<https://github.com/TomMonks/swast-benchmarking>). The computational analyses were run on Intel i9-9900K CPU with 64GB RAM running the Pop!_OS 20.04 Linux. Our benchmark model has been released as a MIT licensed (free and open) Python package. A cloud runnable tutorial, via BinderHub, is available from <https://github.com/TomMonks/swast-forecast-tool>.

2.5 Candidate Models

We evaluated 14 candidate forecasting methods, listed in Table 1, relative to a naive benchmark. We selected established candidate methods (many of which are known to the UK National Health Service); for example, Holt-Winters Exponential Smoothing, Autoregressive Integrated Moving Average (ARIMA), Harmonic regression (regression with ARIMA errors using fourier series exogenous variables to represent seasonality), TBATS (Trigonometric seasonality, Box-Cox transformation, ARMA errors, Trend and Seasonal components) [19], and Singular Spectrum Analysis (SSA). We also included four ensembles (an unweighted average) of methods: three of individual methods already included as well as a standard ensemble *comb* [20].

We use a modern approach to ARIMA model selection by automatic selection of the model using the Hyndman-Khandakar algorithm. The number of fourier terms in harmonic regression were selected by minimising Akaike Information Criterion. The most modern method we employ is Facebook Prophet which was designed to handle higher frequency data that may have multiple periodicity. Prophet is similar to a Generalised Additive Model in that it is a curve fitting approach. All methods are simple for ambulance services to implement and available in either Python or R.

Table 1: **Candidate forecasting methods**

Method	Description
1	Holt-Winters Exponential Smoothing with additive errors and seasonality
2	Automatic Autoregressive Integrated Moving Average (autoARIMA)
3	Lagged regression (autoregression) with holidays and seasonal indexes
4	Lagged regularised regression (elastic-net) with seasonal indexes
5	Regression with holidays and ARIMA errors
6	Regression with holidays, seasonal indexes and ARIMA errors
7	Harmonic Regression (fourier terms) with holidays
8	Singular Spectrum Analysis
9	Facebook Prophet
10	TBATS
11	Comb: Simple Exponential Smoothing, Holt’s Linear Trend, Holt’s damped trend
12	Ensemble of HW [1] and AutoARIMA[2]
13	Ensemble Prophet [7], and Regression with ARIMA Errors [5]
14	Ensemble of HW [1], Prophet [9] and Regression with ARIMA errors [5]

2.6 Statistical analysis

2.6.1 A naive forecasting benchmark

A naive baseline forecasting method was chosen. This was to ensure that the sophisticated methods we test in the study were only considered for the final benchmark if they provided more accurate point forecasts than the simplest of models. As emergency care demand data are seasonal we opted for the well-known Seasonal Naive method [8]. This method works by using the most recent observation for the same day and carrying it forward. For example, if we are forecasting next Tuesday then the observation from the most recent Tuesday is used as the predicted value.

2.6.2 Model selection

Model selection was conducted using a method of time series cross-validation called rolling forecast origin [21]. To avoid leakage of future observations, the method incrementally moves the forecast origin forward in time and then makes a prediction. For each new fold we implemented a stride of seven days. We performed a two stage model selection procedure. In the first stage, we used an aggregate regional level time series to screen and identify the most promising candidate models for up to a 365 day forecast (27 folds). In the second stage, the best models were compared at the regional and seven sub-regional levels. For the 84 day horizon we had sufficient data to produce 67 validation folds for each sub-region.

2.6.3 Simulated forecast setting

Our test set provided 365 observations for each of the seven sub-regions. This enabled us to make 41 simulated forecasts of 7 to 84 days. Our benchmark therefore includes a range of MASE and coverage metrics that might be expected in practice.

2.6.4 External validation

For external validation, we repeated the simulated forecast procedure within 13 test sets for London, Yorkshire and Wales. We tested if the selected model performed as expected relative to a simpler baseline and compared the range of MASE and coverage metrics observed to those expected in our benchmark analysis.

3 Results

3.1 Training data

There were no missing observations in the training data. Figure 1 depicts the training data time series. The median number of calls requiring an ambulance dispatch per day was 2169 (IQR 2083 - 2269). Extreme days were observed on the 1st January every year (new year’s day) where the median number of calls

increased to 2783 (IQR 2673 - 2930). The data displayed a slight upward trend over time rising from median per day of 2135 in 2014 to 2257 by 2018. There was variation in demand by month of year and day of the week (Figure 2).

Figure 1: **Time series of training data.** Trust level daily number of calls that require one or more ambulance dispatches. Extreme observations observed on new years day marked with red dot.

Figure 2: **Annual and weekly seasonality.** Top figure illustrates the variation in demand by month of the year. Bottom figure illustrates variation by day of week.

3.2 Naive benchmark

The magnitude of the point forecasts of the seasonal naive method increased with the forecast horizon (see Table 2). On average, the naive method achieved a MASE of 0.94 (0.35) over 7 days. For the trust, this represents a sMAPE and RMSE of 3.5% and 96.9 calls respectively. Over an 84 day horizon MASE increased to 1.34, (a 30% increase). By 365 days MASE had increased to 1.49 (0.46).

	MASE	sMAPE	RMSE
7	0.94 (0.35)	3.48 (1.19)	96.91 (37.11)
14	1.06 (0.39)	3.94 (1.32)	109.71 (41.88)
21	1.12 (0.41)	4.12 (1.38)	115.38 (42.94)
28	1.14 (0.40)	4.19 (1.33)	118.34 (42.92)
35	1.18 (0.40)	4.35 (1.34)	123.26 (42.33)
42	1.23 (0.40)	4.51 (1.33)	128.11 (40.34)
49	1.25 (0.39)	4.61 (1.32)	131.40 (39.79)
56	1.27 (0.39)	4.67 (1.33)	133.95 (40.20)
63	1.30 (0.40)	4.77 (1.36)	136.59 (40.30)
70	1.31 (0.37)	4.82 (1.26)	138.70 (37.25)
77	1.33 (0.34)	4.88 (1.18)	140.73 (34.49)
84	1.34 (0.35)	4.94 (1.21)	143.04 (34.07)
365	1.49 (0.46)	5.51 (1.64)	155.29 (38.64)

Table 2: **Cross-Validation of Seasonal Naive Point Forecasts.** Figures are mean (standard deviation) Mean Absolute Scaled Error, symmetric Mean Absolute Percentage Error and Root Mean Squared Error by forecast horizon ($n = 27$ folds)

3.3 Model selection

The stage one MASE, 80 and 95% prediction interval coverage cross-validation results are summarised in tables S1, S2 and S3 in Appendix A, respectively. In the first stage of model selection, only Prophet (model 7) and the Prophet-Regression with ARIMA errors ensemble (model 13) had a MASE lower than 1.0 up to 84 day horizon ($MASE_{84}$ Prophet = 0.97 (0.12); ensemble = 0.99 (0.13)). In the ensemble the ARIMA model selected was a (1, 1, 3)(1, 0, 1, 7). The two models out performed seasonal naive at all horizons. All models had a MASE of greater than 1.0 at 365 days.

In the second stage of model selection, we compared Prophet and the ensemble at the sub region level. Figure 3 illustrates the change in MASE by forecast horizon and sub-region. Figure 4 illustrates the distribution of prediction interval coverage by horizon and desired coverage. The MASE for Prophet and the ensemble models is similar across all sub-regions. The median prediction interval coverage for the ensemble is more consistent than Prophet. The upper quartile of Prophet’s prediction interval coverage fails to achieve desired coverage from a horizon of 21 days. We chose the ensemble as our forecast benchmark model.

3.4 Benchmark accuracy in the test set

Across all sub-regions and a horizon of 7 to 84 days the ensemble scored a benchmark MASE of 0.68 (95% CI 0.67 - 0.69) with 90% of the mean point forecast errors between 0.49 and 0.91 (an equivalent average error measured by sMAPE is 4.9%; 95% CI 4.7 - 5.1). Mean coverage for 80 and 95% prediction intervals was 0.847 (95% CI 0.843 - 0.851) and 0.965 (95% CI 0.949 - 0.977), respectively. Table 3 reports MASE and coverage for each forecast horizon. Table 4 provides detailed results of the average coverage for the 60th - 95th prediction intervals for each region.

3.5 External validation

There were no missing data in the external validation set. Overall the ensemble produced a MASE of 0.73 (95% CI 0.72 - 0.74) with 90% of validation folds achieving a MASE between 0.64 - 0.83). The ensembles’

h	MASE	Coverage 80%	Coverage 95%
7	0.66 (0.64 - 0.69)	0.846 (0.829 - 0.864)	0.961 (0.952 - 0.970)
14	0.67 (0.65 - 0.69)	0.847 (0.835 - 0.860)	0.964 (0.958 - 0.970)
21	0.67 (0.65 - 0.68)	0.850 (0.839 - 0.860)	0.964 (0.959 - 0.970)
28	0.67 (0.66 - 0.69)	0.849 (0.840 - 0.858)	0.965 (0.961 - 0.969)
35	0.67 (0.66 - 0.69)	0.849 (0.841 - 0.858)	0.965 (0.961 - 0.969)
42	0.68 (0.66 - 0.69)	0.847 (0.839 - 0.855)	0.965 (0.962 - 0.969)
49	0.68 (0.67 - 0.69)	0.847 (0.839 - 0.855)	0.965 (0.962 - 0.969)
56	0.68 (0.67 - 0.70)	0.847 (0.839 - 0.854)	0.966 (0.963 - 0.969)
63	0.69 (0.68 - 0.70)	0.846 (0.839 - 0.854)	0.966 (0.964 - 0.969)
70	0.69 (0.68 - 0.70)	0.845 (0.838 - 0.853)	0.967 (0.964 - 0.969)
77	0.70 (0.68 - 0.71)	0.844 (0.837 - 0.852)	0.967 (0.964 - 0.969)
84	0.70 (0.69 - 0.71)	0.843 (0.835 - 0.851)	0.966 (0.963 - 0.969)

Table 3: **Benchmark results by forecast horizon.** Mean (95% CI) MASE and 80, 95% prediction interval coverage. Figures pool all simulated forecasts from all seven sub-regions ($n = 287$; 41 folds per sub-region).

	60%	70%	80%	90%	95%
BNSSG	0.639 (0.634 - 0.644)	0.746 (0.739 - 0.753)	0.844 (0.839 - 0.850)	0.922 (0.918 - 0.926)	0.959 (0.954 - 0.963)
Cornwall	0.637 (0.634 - 0.639)	0.762 (0.760 - 0.764)	0.856 (0.855 - 0.857)	0.921 (0.921 - 0.922)	0.974 (0.974 - 0.975)
Devon	0.691 (0.684 - 0.699)	0.793 (0.789 - 0.797)	0.857 (0.856 - 0.859)	0.939 (0.935 - 0.942)	0.974 (0.972 - 0.977)
Dorset	0.712 (0.708 - 0.715)	0.803 (0.800 - 0.805)	0.873 (0.871 - 0.876)	0.940 (0.938 - 0.941)	0.975 (0.974 - 0.976)
Gloucestershire	0.626 (0.621 - 0.632)	0.730 (0.728 - 0.733)	0.835 (0.833 - 0.837)	0.936 (0.934 - 0.937)	0.963 (0.962 - 0.964)
Somerset	0.629 (0.625 - 0.633)	0.723 (0.718 - 0.728)	0.809 (0.805 - 0.813)	0.901 (0.899 - 0.903)	0.950 (0.948 - 0.951)
Wiltshire	0.665 (0.662 - 0.668)	0.762 (0.760 - 0.765)	0.853 (0.847 - 0.859)	0.928 (0.925 - 0.932)	0.961 (0.960 - 0.963)

Table 4: **Prediction interval coverage by sub-region.** Mean (95% CI) prediction interval coverage. Figures relate to 41 folds of 7-84 days in the test set.

Figure 3: **Cross-validation variation in MASE across regions by horizon.** Shaded area is 95% confidence intervals for the mean point forecast error of the Prophet and Ensemble models.

80% and 95% prediction intervals provided 0.833 (95% CI 0.828-0.838) and 0.965 95% CI (0.963-0.967) coverage, respectively. Table 5 reports MASE and coverage by region.

4 Discussion

Our chosen benchmark method, based on performance, is an ensemble (a simple average) of Facebook’s Prophet and Regression with ARIMA errors. Both methods are flexible enough to add in special calendar events such as national holidays. In our model we chose to include new year’s day as this clearly stood out in the time series. In our regression model, we model the error process using the same ARIMA model - $(1, 1, 3)(1, 0, 1, 7)$ - for each sub region. EMS providers in different regions may wish to experiment with alternative error processes.

Our cross-validation demonstrated that performance of the ensemble was superior to either method on its own, the other candidate models and a naive benchmark. However, we note that Prophet is also a reasonable choice for ambulance trusts new to forecasting (albeit they should recognise the shortcomings in terms of coverage). We emphasise the critical importance of a naive benchmark such as seasonal naive in cross-validation to confirm that more complex models add value. We found that over our forecast horizon seasonal naive outperformed several state-of-the-art forecasting techniques.

We encourage forecasters in the ambulance service to use both point forecasts and prediction intervals. A singular focus on a point forecast is unwise; it is not possible to predict the future exactly and so forecasters should take account of the range of likely values. We found that the most accurate single method for point forecasts did not produce satisfactory coverage. The software we have developed to support forecasting in ambulance services reports 95% prediction intervals by default.

Figure 4: **Cross-validation variation in Prediction Interval Coverage across regions by horizon.** The box plots illustrate the distribution of coverage across all six regions. Red horizontal lines represent desired coverage. Top and bottom rows represents 95% and 80% prediction intervals, respectively. Left and right columns represent Prophet and the Ensemble, respectively. It is desirable to achieve coverage, but not exceed substantially.

Turning to benchmark performance, our simulated forecast achieved a MASE of 0.68. We found this performance declined slightly over the forecast horizon, and 90% of forecasts fell into range of 0.49 and 0.91. The latter is a reasonable approximation for practitioners to use as a rule of thumb for benchmarking. Researchers should make use of our detailed breakdown of MASE to enable simple robust comparison across regions and studies. We emphasise that in ‘real terms’ for an individual ambulance service our forecasts demonstrate a useful accuracy up to 84 days. For instance, the RMSE for a 84 day horizon in Cornwall (average calls per day = 243) and Dorset (average = 323) was between 16 and 18 calls.

Our ensemble method achieved desired coverage, but we acknowledge that it is conservative at the 80% level. We found that coverage varied by sub-region and advise practitioners to investigate their own time series. A benchmark coverage for 84% and 96% was achieved for the 80 and 95% prediction intervals, respectively. We provide a more detailed breakdown of coverage in table 4 for future scientific forecasting studies. Our detailed results are also relevant to studies that aim to building forecasting into decision support systems [22].

There are several strengths to our study. First, our model selection focused on well-known forecasting methods including a recognised and relevant naive benchmark. We excluded novel methods development from our study as we wished our method to be widely available and easily implemented in a health service - particularly the UK’s NHS where the study setting is based. Indeed the methods in our ensemble are recognised by NHS Improvement for forecasting [23]. Our method is implemented and available in Python and can easily be implemented in R. Second, we based our benchmark accuracy on the performance of the ensemble on seven time series from regions in South West of England. The results were replicated in 13 regions from London, Wales and Yorkshire. This evidence provides a robust estimate of what can be expected in practice and a strong comparator for future studies attempting to improve on our ensemble. Third, we include prediction interval coverage as a primary outcome in addition to point forecast accuracy. We note that prediction coverage is a critical measure of the performance of a forecast and for making a robust decision, but typically forecasting studies ignore it.

Trust	Region	MASE	Coverage 80	Coverage 95
LAS	North Central	0.76 (0.11)	0.84 (0.06)	0.94 (0.04)
	North East	0.74 (0.08)	0.87 (0.06)	0.97 (0.03)
	North West	0.75 (0.11)	0.86 (0.07)	0.97 (0.03)
	South East	0.70 (0.10)	0.85 (0.06)	0.98 (0.03)
	South West	0.68 (0.10)	0.86 (0.06)	0.97 (0.04)
WAST	Central and West	0.69 (0.12)	0.82 (0.09)	0.96 (0.04)
	North	0.73 (0.14)	0.79 (0.10)	0.97 (0.04)
	South East	0.64 (0.10)	0.87 (0.06)	0.98 (0.04)
YAS	ABL	0.67 (0.11)	0.86 (0.05)	0.98 (0.02)
	CKW	0.74 (0.09)	0.82 (0.07)	0.96 (0.03)
	Humb and ER	0.77 (0.11)	0.79 (0.08)	0.94 (0.04)
	North Yorks	0.83 (0.15)	0.77 (0.09)	0.94 (0.06)
	South	0.76 (0.11)	0.81 (0.08)	0.96 (0.04)

Table 5: **External validation: Point forecast and coverage performance by region.** Mean (SD). Figures calculated from 41 folds of 7-84 days in test sets.

Our study also has several limitations. Our geographic regions are limited in number and within England and Wales only. Other geographic regions of the UK and other nations may have differing seasonal patterns. We argue that two elements of our study mitigate this issue. First our approach makes use of automatic modelling procedures that enable deployment at scale and manage some issues with differences in seasonality. Second, our primary objective is clear benchmark performance. Other regions can easily compare their outcomes to our own. Our data run until January 2020 and takes place prior to the COVID-19 pandemic. In this time period emergency services around the world will have seen dramatic shifts in their demand and a potential increase in the importance of weather. Assuming demand returns to similar patterns post-pandemic our ensemble should be viable with some minor modification. For instance, both Prophet and Regression with ARIMA errors can be modified to include binary ‘intervention’ variables that can represent different phases of the pandemic. Prophet’s flexible modelling of trend also allows for manual correction at critical time points (e.g. the start and end of lockdowns). A further limitation is that our methods do not make use of other features such as weather forecasts.

4.1 Conclusions

The primary contribution of our study is the benchmark performance for predicting the number of calls that result in the dispatch of one or more ambulances. We provide estimates up to 84 days in advance. Future studies and novel methodologies should now aim to exceed these benchmarks. An exciting future direction is the study of machine learning methods including neural networks to investigate if they can match the performance of our benchmark. For methods to be accepted as state-of-the-art they must now provide evidence that their performance exceeds both a naive method, and the point forecasts and coverage benchmarks generated by our ensemble. We particularly encourage researchers to improve on our prediction interval coverage.

References

- [1] D. W. Lee, H. J. Moon, and N. H. Heo. “Association between ambulance response time and neurologic outcome in patients with cardiac arrest”. en. In: *The American Journal of Emergency Medicine* 37.11 (Nov. 2019), pp. 1999–2003. ISSN: 0735-6757. DOI: 10.1016/j.ajem.2019.02.021. URL: <http://www.sciencedirect.com/science/article/pii/S0735675719300993> (visited on 07/27/2020).
- [2] A. Meretoja, M. Keshtkaran, J. L. Saver, et al. “Stroke thrombolysis: Save a minute, save a day”. In: *Stroke* 45.4 (2014), pp. 1053–1058. ISSN: 15244628. DOI: 10.1161/STROKEAHA.113.002910.
- [3] J. P. Byrne, N. C. Mann, M. Dai, et al. “Association Between Emergency Medical Service Response Time and Motor Vehicle Crash Mortality in the United States”. en. In: *JAMA Surgery* 154.4 (Apr. 2019). Publisher: American Medical Association, pp. 286–293. ISSN: 2168-6254. DOI: 10.1001/jamasurg.2018.5097. URL: <https://jamanetwork.com/journals/jamasurgery/fullarticle/2723267> (visited on 07/27/2020).

- [4] J. R. Baker and K. E. Fitzpatrick. “Determination of an Optimal Forecast Model for Ambulance Demand Using Goal Programming”. In: *The Journal of the Operational Research Society* 37.11 (1986), pp. 1047–1059. ISSN: 01605682, 14769360. URL: <http://www.jstor.org/stable/2582689>.
- [5] J. L. Vile, J. W. Gillard, P. R. Harper, et al. “Predicting ambulance demand using singular spectrum analysis”. In: *Journal of the Operational Research Society* 63.11 (2012), pp. 1556–1565. DOI: 10.1057/jors.2011.160. URL: <https://doi.org/10.1057/jors.2011.160>.
- [6] H. Setzler, C. Saydam, and S. Park. “EMS call volume predictions: A comparative study”. In: *Computers Operations Research* 36.6 (2009), pp. 1843–1851. ISSN: 0305-0548. DOI: <https://doi.org/10.1016/j.cor.2008.05.010>. URL: <http://www.sciencedirect.com/science/article/pii/S0305054808001068>.
- [7] L. H. Brown, E. B. Lerner, B. Larmon, et al. “Are EMS Call Volume Predictions Based on Demand Pattern Analysis Accurate?” In: *Prehospital Emergency Care* 11.2 (2007), pp. 199–203. DOI: 10.1080/10903120701204797. URL: <https://doi.org/10.1080/10903120701204797>.
- [8] R. Hyndman and G. Athanasopoulos. *Forecasting: principles and practice, 2nd edition*. Online accessed on [current date]. OTexts: Melbourne, Australia., 2018. URL: OTexts.com/fpp2.
- [9] R. J. Hyndman and A. B. Koehler. “Another look at measures of forecast accuracy”. In: *International Journal of Forecasting* 22.4 (2006), pp. 679–688. ISSN: 0169-2070. DOI: <https://doi.org/10.1016/j.ijforecast.2006.03.001>. URL: <http://www.sciencedirect.com/science/article/pii/S0169207006000239>.
- [10] T. G. Smith et al. *pmdarima: ARIMA estimators for Python*. [Online; accessed [today]]. 2017. URL: <http://www.alkaline-ml.com/pmdarima>.
- [11] S. Taylor and B. Letham. “Forecasting at scale”. In: *PeerJ Preprints* (2017). DOI: 5:e3190v2. URL: <https://doi.org/10.7287/peerj.preprints.3190v2>.
- [12] S. Seabold and J. Perktold. “statsmodels: Econometric and statistical modeling with python”. In: *9th Python in Science Conference*. 2010.
- [13] G. Skorupa. *BATS and TBATS time series forecasting*. [Online; accessed [today]]. 2020. URL: <https://pypi.org/project/tbats/>.
- [14] T. Monks. *forecast-tools: fundamental tools to support the forecasting process in python*. 2020. DOI: 10.5281/zenodo.3969789. URL: <https://doi.org/10.5281/zenodo.3969789>.
- [15] A. Korobeynikov. “Computation- and space-efficient implementation of SSA”. In: *Statistics and Its Interface* 3.3 (2010). R package version 1.0.2, pp. 357–368.
- [16] W. McKinney. “pandas: a foundational Python library for data analysis and statistics”. In: *Python for High Performance and Scientific Computing* 14 (2011).
- [17] S. van der Walt, S. C. Colbert, and G. Varoquaux. “The NumPy Array: A Structure for Efficient Numerical Computation”. In: *Computing in Science Engineering* 13.2 (2011), pp. 22–30. DOI: 10.1109/MCSE.2011.37.
- [18] J. D. Hunter. “Matplotlib: A 2D graphics environment”. In: *Computing in Science & Engineering* 9.3 (2007), pp. 90–95. DOI: 10.1109/MCSE.2007.55.
- [19] A. M. D. Livera, R. J. Hyndman, and R. D. Snyder. “Forecasting Time Series With Complex Seasonal Patterns Using Exponential Smoothing”. In: *Journal of the American Statistical Association* 106.496 (2011), pp. 1513–1527. DOI: 10.1198/jasa.2011.tm09771. URL: <https://doi.org/10.1198/jasa.2011.tm09771>.
- [20] S. Makridakis, E. Spiliotis, and V. Assimakopoulos. “The M4 Competition: 100,000 time series and 61 forecasting methods”. In: *International Journal of Forecasting* 36.1 (2020). M4 Competition, pp. 54–74. ISSN: 0169-2070. DOI: <https://doi.org/10.1016/j.ijforecast.2019.04.014>. URL: <http://www.sciencedirect.com/science/article/pii/S0169207019301128>.
- [21] L. J. Tashman. “Out-of-sample tests of forecasting accuracy: an analysis and review”. In: *International Journal of Forecasting* 16.4 (2000). The M3- Competition, pp. 437–450. ISSN: 0169-2070. DOI: [https://doi.org/10.1016/S0169-2070\(00\)00065-0](https://doi.org/10.1016/S0169-2070(00)00065-0). URL: <http://www.sciencedirect.com/science/article/pii/S0169207000000650>.
- [22] J. Vile, J. Gillard, P. Harper, et al. “Time-dependent stochastic methods for managing and scheduling Emergency Medical Services”. In: *Operations Research for Health Care* 8 (2016), pp. 42–52. ISSN: 2211-6923. DOI: <https://doi.org/10.1016/j.orhc.2015.07.002>. URL: <http://www.sciencedirect.com/science/article/pii/S2211692314200518>.
- [23] NHS England. *NHS England Advanced forecasting techniques*. Tech. rep. 2020. URL: <https://www.england.nhs.uk/publication/advanced-forecasting-techniques/> (visited on 03/17/2021).

Appendices

A Stage 1 cross-validation

h	naive	1	2	3	4	5	6	7	8	9	10	11	12	13	14
7	0.94 (0.35)	0.78 (0.37)	0.82 (0.39)	0.86 (0.32)	0.76 (0.28)	1.21 (0.67)	0.88 (0.34)	0.88 (0.46)	1.01 (0.46)	0.81 (0.42)	0.85 (0.59)	1.12 (0.45)	0.77 (0.33)	0.78 (0.34)	0.77 (0.35)
14	1.06 (0.39)	0.88 (0.43)	0.92 (0.40)	0.96 (0.31)	0.85 (0.29)	1.62 (0.90)	0.97 (0.31)	0.97 (0.44)	1.07 (0.42)	0.82 (0.36)	0.95 (0.63)	1.19 (0.44)	0.87 (0.36)	0.83 (0.32)	0.84 (0.35)
21	1.12 (0.41)	0.92 (0.42)	0.97 (0.37)	1.02 (0.27)	0.91 (0.29)	1.88 (1.02)	1.03 (0.26)	1.03 (0.38)	1.13 (0.40)	0.83 (0.29)	1.01 (0.68)	1.22 (0.41)	0.91 (0.35)	0.85 (0.26)	0.86 (0.30)
28	1.14 (0.40)	0.95 (0.39)	1.01 (0.36)	1.07 (0.27)	0.96 (0.28)	2.08 (1.09)	1.09 (0.25)	1.07 (0.35)	1.16 (0.39)	0.83 (0.23)	1.05 (0.68)	1.24 (0.37)	0.94 (0.33)	0.86 (0.20)	0.88 (0.25)
35	1.18 (0.40)	1.00 (0.40)	1.05 (0.37)	1.12 (0.29)	1.00 (0.29)	2.27 (1.17)	1.14 (0.27)	1.12 (0.41)	1.19 (0.40)	0.84 (0.20)	1.10 (0.69)	1.28 (0.37)	0.98 (0.34)	0.88 (0.17)	0.91 (0.23)
42	1.23 (0.40)	1.04 (0.40)	1.08 (0.37)	1.15 (0.31)	1.05 (0.31)	2.44 (1.23)	1.17 (0.27)	1.17 (0.43)	1.21 (0.40)	0.84 (0.18)	1.15 (0.69)	1.31 (0.38)	1.02 (0.34)	0.90 (0.17)	0.93 (0.23)
49	1.25 (0.39)	1.08 (0.40)	1.11 (0.36)	1.19 (0.32)	1.08 (0.31)	2.60 (1.22)	1.20 (0.27)	1.20 (0.43)	1.23 (0.40)	0.85 (0.16)	1.18 (0.69)	1.34 (0.36)	1.05 (0.35)	0.91 (0.16)	0.95 (0.22)
56	1.27 (0.39)	1.10 (0.38)	1.13 (0.34)	1.22 (0.32)	1.11 (0.30)	2.74 (1.16)	1.23 (0.27)	1.24 (0.43)	1.25 (0.40)	0.86 (0.15)	1.19 (0.68)	1.36 (0.34)	1.07 (0.34)	0.93 (0.14)	0.97 (0.21)
63	1.30 (0.40)	1.13 (0.39)	1.15 (0.33)	1.25 (0.33)	1.14 (0.31)	2.87 (1.10)	1.25 (0.26)	1.27 (0.43)	1.27 (0.38)	0.87 (0.13)	1.22 (0.66)	1.39 (0.34)	1.10 (0.33)	0.95 (0.13)	0.99 (0.21)
70	1.31 (0.37)	1.15 (0.36)	1.17 (0.31)	1.28 (0.33)	1.17 (0.31)	2.97 (1.01)	1.27 (0.26)	1.30 (0.44)	1.29 (0.36)	0.87 (0.12)	1.23 (0.65)	1.42 (0.32)	1.12 (0.30)	0.96 (0.12)	1.01 (0.19)
77	1.33 (0.34)	1.17 (0.33)	1.18 (0.29)	1.31 (0.33)	1.19 (0.31)	3.07 (0.90)	1.29 (0.26)	1.33 (0.44)	1.31 (0.34)	0.88 (0.12)	1.24 (0.62)	1.44 (0.29)	1.14 (0.28)	0.97 (0.11)	1.02 (0.17)
84	1.34 (0.35)	1.19 (0.33)	1.20 (0.30)	1.34 (0.33)	1.21 (0.33)	3.18 (0.84)	1.32 (0.26)	1.37 (0.44)	1.33 (0.33)	0.90 (0.11)	1.25 (0.60)	1.47 (0.29)	1.16 (0.29)	0.99 (0.13)	1.04 (0.18)
365	1.49 (0.46)	1.39 (0.46)	1.34 (0.41)	1.33 (0.06)	1.52 (0.59)	3.05 (0.95)	1.33 (0.07)	1.60 (0.41)	1.44 (0.12)	1.15 (0.20)	1.37 (0.49)	1.65 (0.47)	1.39 (0.47)	1.26 (0.33)	1.29 (0.37)

Table S1: Results for stage 1 model selection. Mean Absolute Scaled Error (MASE) of Point Forecasts (Stdev). The numeric numbering of columns references the list of methods in Table 1.

1	2	3	4	5	6	7	8	9	10	11	12	13	14
7	0.85 (0.19)	0.81 (0.20)	0.67 (0.21)	0.80 (0.18)	0.78 (0.29)	0.68 (0.23)	0.80 (0.23)	0.78 (0.20)	0.84 (0.23)	0.84 (0.14)	0.83 (0.21)	0.79 (0.21)	0.82 (0.20)
14	0.86 (0.18)	0.78 (0.17)	0.61 (0.18)	0.79 (0.16)	0.76 (0.28)	0.61 (0.19)	0.78 (0.20)	0.77 (0.17)	0.85 (0.23)	0.91 (0.10)	0.82 (0.19)	0.78 (0.18)	0.81 (0.19)
21	0.87 (0.16)	0.77 (0.15)	0.57 (0.15)	0.77 (0.15)	0.74 (0.27)	0.57 (0.16)	0.78 (0.15)	0.76 (0.14)	0.86 (0.22)	0.94 (0.06)	0.83 (0.17)	0.78 (0.14)	0.82 (0.16)
28	0.88 (0.14)	0.76 (0.14)	0.54 (0.14)	0.74 (0.14)	0.71 (0.27)	0.54 (0.14)	0.76 (0.14)	0.66 (0.25)	0.76 (0.12)	0.87 (0.21)	0.84 (0.14)	0.78 (0.11)	0.83 (0.13)
35	0.89 (0.14)	0.75 (0.15)	0.52 (0.15)	0.72 (0.14)	0.69 (0.28)	0.52 (0.15)	0.75 (0.16)	0.66 (0.26)	0.77 (0.11)	0.88 (0.21)	0.84 (0.14)	0.77 (0.10)	0.83 (0.13)
42	0.90 (0.14)	0.74 (0.15)	0.50 (0.15)	0.71 (0.14)	0.66 (0.28)	0.50 (0.16)	0.73 (0.16)	0.66 (0.26)	0.76 (0.10)	0.89 (0.20)	0.84 (0.14)	0.77 (0.10)	0.84 (0.12)
49	0.90 (0.12)	0.73 (0.15)	0.49 (0.15)	0.70 (0.14)	0.64 (0.29)	0.49 (0.15)	0.72 (0.17)	0.66 (0.25)	0.76 (0.10)	0.89 (0.20)	0.85 (0.13)	0.76 (0.09)	0.84 (0.11)
56	0.91 (0.10)	0.73 (0.14)	0.47 (0.14)	0.69 (0.14)	0.63 (0.28)	0.47 (0.15)	0.71 (0.18)	0.65 (0.24)	0.76 (0.10)	0.90 (0.19)	0.85 (0.12)	0.76 (0.08)	0.84 (0.10)
63	0.91 (0.10)	0.72 (0.14)	0.45 (0.14)	0.69 (0.15)	0.62 (0.26)	0.46 (0.14)	0.70 (0.18)	0.65 (0.24)	0.76 (0.09)	0.90 (0.18)	0.86 (0.12)	0.75 (0.08)	0.84 (0.09)
70	0.92 (0.09)	0.72 (0.13)	0.44 (0.14)	0.68 (0.15)	0.61 (0.23)	0.45 (0.14)	0.69 (0.18)	0.64 (0.23)	0.76 (0.09)	0.91 (0.18)	0.86 (0.11)	0.75 (0.07)	0.85 (0.09)
77	0.92 (0.09)	0.72 (0.12)	0.43 (0.14)	0.68 (0.15)	0.60 (0.21)	0.43 (0.13)	0.69 (0.18)	0.63 (0.22)	0.75 (0.09)	0.91 (0.17)	0.86 (0.10)	0.74 (0.07)	0.85 (0.08)
84	0.93 (0.08)	0.71 (0.12)	0.41 (0.13)	0.67 (0.16)	0.58 (0.20)	0.42 (0.12)	0.67 (0.18)	0.62 (0.21)	0.75 (0.09)	0.92 (0.16)	0.87 (0.10)	0.73 (0.09)	0.85 (0.08)
365	0.97 (0.02)	0.78 (0.11)	0.43 (0.04)	0.73 (0.16)	0.79 (0.12)	0.45 (0.04)	0.63 (0.19)	0.59 (0.11)	0.67 (0.10)	0.98 (0.03)	0.94 (0.06)	0.72 (0.13)	0.91 (0.07)

Table S2: Cross Validation of 80% Prediction Interval Coverage (cv=27)

	1	2	3	4	5	6	7	8	9	10	11	12	13	14
7	0.97 (0.07)	0.96 (0.09)	0.90 (0.15)	0.96 (0.08)	0.93 (0.16)	0.89 (0.16)	0.94 (0.15)	0.90 (0.19)	0.92 (0.14)	0.97 (0.06)	0.96 (0.06)	0.97 (0.08)	0.94 (0.13)	0.94 (0.13)
14	0.96 (0.11)	0.95 (0.11)	0.86 (0.16)	0.94 (0.11)	0.91 (0.18)	0.85 (0.17)	0.94 (0.12)	0.91 (0.16)	0.91 (0.11)	0.97 (0.08)	0.97 (0.08)	0.95 (0.11)	0.93 (0.11)	0.92 (0.14)
21	0.96 (0.10)	0.95 (0.10)	0.83 (0.15)	0.93 (0.11)	0.90 (0.19)	0.82 (0.15)	0.94 (0.10)	0.91 (0.15)	0.91 (0.09)	0.97 (0.06)	0.98 (0.04)	0.95 (0.11)	0.92 (0.10)	0.93 (0.11)
28	0.97 (0.08)	0.94 (0.09)	0.80 (0.14)	0.93 (0.10)	0.89 (0.20)	0.79 (0.13)	0.94 (0.09)	0.91 (0.14)	0.91 (0.08)	0.98 (0.05)	0.99 (0.03)	0.95 (0.09)	0.92 (0.08)	0.94 (0.09)
35	0.97 (0.07)	0.94 (0.09)	0.78 (0.14)	0.92 (0.10)	0.89 (0.19)	0.77 (0.14)	0.93 (0.09)	0.91 (0.14)	0.91 (0.07)	0.98 (0.04)	0.99 (0.02)	0.95 (0.08)	0.92 (0.07)	0.94 (0.08)
42	0.98 (0.05)	0.94 (0.08)	0.75 (0.15)	0.91 (0.10)	0.88 (0.19)	0.75 (0.14)	0.92 (0.09)	0.91 (0.14)	0.91 (0.07)	0.98 (0.03)	0.99 (0.02)	0.96 (0.07)	0.92 (0.06)	0.94 (0.07)
49	0.98 (0.05)	0.93 (0.08)	0.73 (0.16)	0.91 (0.09)	0.87 (0.19)	0.74 (0.14)	0.91 (0.09)	0.91 (0.14)	0.91 (0.06)	0.98 (0.03)	0.99 (0.02)	0.96 (0.06)	0.92 (0.06)	0.95 (0.06)
56	0.98 (0.04)	0.93 (0.08)	0.71 (0.16)	0.90 (0.08)	0.88 (0.16)	0.72 (0.13)	0.90 (0.10)	0.91 (0.14)	0.91 (0.06)	0.98 (0.02)	0.99 (0.01)	0.96 (0.06)	0.92 (0.05)	0.95 (0.06)
63	0.98 (0.04)	0.93 (0.07)	0.70 (0.16)	0.89 (0.09)	0.87 (0.16)	0.71 (0.13)	0.89 (0.11)	0.91 (0.12)	0.90 (0.06)	0.99 (0.02)	0.99 (0.01)	0.96 (0.05)	0.92 (0.04)	0.95 (0.05)
70	0.98 (0.03)	0.92 (0.06)	0.68 (0.16)	0.89 (0.09)	0.87 (0.15)	0.70 (0.13)	0.88 (0.12)	0.91 (0.11)	0.90 (0.05)	0.99 (0.02)	0.99 (0.01)	0.96 (0.05)	0.92 (0.04)	0.95 (0.05)
77	0.98 (0.03)	0.92 (0.06)	0.67 (0.16)	0.88 (0.09)	0.87 (0.13)	0.69 (0.13)	0.88 (0.13)	0.91 (0.10)	0.90 (0.05)	0.99 (0.02)	0.99 (0.01)	0.96 (0.04)	0.92 (0.03)	0.95 (0.04)
84	0.98 (0.03)	0.92 (0.06)	0.66 (0.16)	0.88 (0.10)	0.87 (0.13)	0.68 (0.12)	0.87 (0.14)	0.91 (0.09)	0.90 (0.05)	0.99 (0.02)	0.99 (0.01)	0.96 (0.04)	0.91 (0.03)	0.95 (0.04)
365	0.99 (0.01)	0.94 (0.04)	0.68 (0.03)	0.91 (0.06)	0.93 (0.06)	0.69 (0.03)	0.84 (0.11)	0.87 (0.05)	0.88 (0.05)	1.00 (0.00)	1.00 (0.00)	0.98 (0.01)	0.92 (0.05)	0.98 (0.01)

Table S3: Cross Validation of 95% Prediction Interval Coverage (cv=27)