Journal of Medical Systems (2019) 43: 288
https://doi.org/10.1007/s10916-019-1418-y

SYSTEMS-LEVEL QUALITY IMPROVEMENT

Predicting Outpatient Appointment Demand Using Machine Learning and Traditional Methods

Brian Klute[1] · Andrew Homb[2] · Wei Chen[1] · Aaron Stelpflug[1]

Received: 21 May 2019 / Accepted: 8 July 2019 / Published online: 19 July 2019

Abstract

Traditional methods have long been used for clinical demand forecasting. Machine learning methods represent the next evolution in forecasting, but model choice and optimization remain challenging for achieving optimal results. To determine the best method to predict demand for outpatient appointments comparing machine learning and traditional methods, this retrospective study analyzed "appointment requests" at a major outpatient department in a destination medical center. Two separate locations (A and B) were assessed with 20 traditional, hybrid (traditional + machine learning) and machine learning methods to determine the best forecasting outcome (lowest Forecast Standard Error, FSE). Data characteristics from both datasets were examined. 20 forecasting models were then assessed and compared for the best result. Location A's data displayed a cyclical and non-trending pattern while Location B's displayed a cyclical and trending pattern. Both Location A and B yielded the feature engineered XGBoost model (machine learning) with the lowest out-of-sample FSE. It is important to carefully analyze and understand the underlying data set pattern and then test a variety of traditional, machine learning, and hybrid prediction methods to achieve optimal predictive results. Additionally, the use of feature engineering or hybrid methods can augment the usefulness of machine learning methods.

Keywords Machine learning · Traditional methods · Outpatient appointment · Forecasting

Background and significance

A full literature search was performed and the context of all resources was used in the research article below [1–29].

The ability to accurately forecast patient appointments at any medical center is crucial in several ways [1]. First, it improves financial health by allowing the organization to plan the appropriate resources in terms of providers and patient slots, avoiding over or under capacity costs [1]. Secondly, it improves patient satisfaction by creating an adaptable scheduling system that avoids long patient delays to be scheduled for an appointment and makes sure that all the necessary resources are available for an optimum patient experience during that appointment. The goal is the right patient with the right provider at the right time with the right resources. Success means matching the supply of providers and other resources (medical device, procedure room, etc.) with patient demand. Failure means inefficient use of those resources or the inability to service patient demand in a satisfactory time period. Particularly in practices where fixed cost capital investment is high, a more accurate expectation of patient demand can help to justify the purchase and facilitate efficient use of advanced equipment, benefiting both the organization and the patients.

The underlying datasets used for forecasting are complex and different for every department in a medical center, exhibiting various trending and cyclical behaviors as well as noise. Forecasting also has to account for both external and internal referral demand which can cause a different cascade of events depending on the initial referral outcome, often outside the initial department. Therefore, accurate forecasting methods need to be adaptable to a variety of demand series patterns and a method that finds the signal in multiple structures is essential.

This article is part of the Topical Collection on *Systems-Level Quality Improvement*

✉ Wei Chen
Chen.Wei@mayo.edu

[1] Department of Management Engineering and Internal Consulting, Mayo Clinic, 200 First Street SW, Rochester, MN 55905, USA

[2] Department of Radiology, Mayo Clinic, Rochester, MN, USA

Table 1 Model Descriptions [2]

Model	Approach	Method
Autoregressive	Uses the dependent relationship between an observation and some number of lagged observations.	Traditional
Box-Cox Transformed Arima	Arima method with power transforms, including, but not limited to, log, square root, and reciprocal transforms of the data.	Traditional
Convolution-Long Short Term Memory Neural Network (LSTM)	A combination of CNNs and LSTMs where the LSTM units read input data using the convolutional process of a CNN.	Machine Learning
Convolutional Neural Network (CNN)	Type of neural network developed for two- dimensional image data, although they can be used for one-dimensional data such as sequences of text and time series.	Machine Learning
Ensemble	Combination of XGBoost and linear regression methods.	Hybrid
Grid Search Arima	Combines the use of autoregressive and moving average methods and the use of differencing of raw observations	Traditional
Grid Search CNN	CNN network with a grid of inputs.	Machine Learning
Grid Search LSTM	LSTM network with a grid of inputs.	Machine Learning
Grid Search MLP Neural Network	MLP network with a grid of inputs.	Machine Learning
Grid Search Naive and Average Forecast	Using the previous observation and averaging prior values.	Traditional
Grid Search Sarima	Arima method for modeling univariate time series data that may contain trend and seasonal components.	Traditional
Grid Search Triple Exponential Smoothing Model (ETS)	An extension of Exponential Smoothing that explicitly adds support for seasonality	Traditional
K-nearest Neighbor	Estimates how likely a data point is to be a member of one group or the other depending on what group the data points nearest to it are in. Classifies new cases based on a similarity measure (e.g., distance functions)	Traditional
Linear Regression	Output value based on a linear combination of input values.	Traditional
LSTM Neural Network	Type of recurrent neural network capable of learning order dependence in sequence prediction problems. Has an internal memory allowing it to accumulate internal state.	Machine Learning
MLP neural network	Feedforward artificial neural network that generates a set of outputs from a set of inputs. An MLP is characterized by several layers of input nodes connected as a directed graph between the input and output layers.	Machine Learning
Moving Average	Uses the dependency between an observation and residual errors from a moving average model applied to lag observations.	Traditional
Naïve Model	Using the previous observation directly as the forecast without any change.	Traditional
Residual Error	Modeling the difference between what was expected and what was predicted is called the residual error	Traditional
XGBoost	Provides a parallel tree boosting (also known as GBDT, GBM) to a gradient boosting method.	Machine Learning

Depending on the structure of the underlying series, an approach that handles both the trends and the cyclical influences will understand the temporal structure well enough to achieve consistently parsimonious results.

Machine learning methods may offer a lot of promise for time series forecasting, specifically the automatic learning of temporal dependence and the automatic handling of temporal structures like trends [2]. Unfortunately, machine learning methods are often developed for univariate time series forecasting problems and often perform worse than classical and even naive forecasting methods leading to their general dismissal as being unsuitable for time series forecasting in general. Nevertheless, deep learning methods can be effective at more complex time series forecasting problems that involve large amounts of data, multiple variates with complex relationships, and even multi-step and time series classification tasks.

Objectives

The primary aim of this study is to compare various machine learning and traditional methods to determine the best method to predict demand for outpatient appointments. The forecasting system designed and tested in this study is intended to find the balance in trending and cyclical influences. Preparation techniques intended to train the underlying structure will aim to reduce the noisy input into the methods. Ultimately, the intent is to reduce bias and variance in the validated forecasts to achieve the lowest Forecast Standard Error (FSE).

Methods

This retrospective study analyzed "appointment requests" at a major outpatient department in a destination medical center at two different locations (A and B). A naïve persistence model was employed to establish a baseline solution. 19 additional time series continuous variable predictions solutions were obtained by training various models using machine learning, traditional, and hybrid methods (Table 1).

Since machine learning models cannot understand temporal data, explicit time-based feature engineering is applied to predict time series and use its pattern learning capabilities to enhance time series predictions. The approach is to choose the best methods that predict important components of a time series, or a combination of the best methods.

This was particularly useful in the tree-based methods like XGBoost (machine learning) which is a space splitting algorithm. Feature engineering is used to transform the inputs of the raw time series into predictable components and the inputs are formulated as a supervised regression to build its predications. The new features combined with use of stationary data demonstrated that a method typically used for event predictions can be used effectively for time series predictions. The additional predictors used are: [3].

- Day of the month
- Day of the week
- Ordinal date (Number of days start of training series)
- Hour of the day

Furthermore, since this space splitting approach is known to work well for series that exhibit cyclical patterns but does not capture trending with the same effectiveness, an ensemble model (hybrid) is created by combining the XGBoost model (machine learning) and a linear regression model (traditional) that is often useful for trending structures [2]. Should a time series exhibit both trending and cyclical behavior, it is postulated that this hybrid method would produce the right balance of bias and variance reduction [2].

Data collection and population characteristics

For location A, 39,205 appointment requests were extracted from that department between August 2016 and August 2018. For location B, 4,538 appointment requests were extracted for that department between September 2016 and September 2017.

For both locations, data was grouped into weekly cycles. Weekly cycles were chosen because they represent the patterns associated with a destination medical center where consultation, testing, diagnosis and follow-up visits are often occurring within one week.

The data obtained was initially examined to determine the data pattern. The raw time series for both locations is plotted as Fig. 1. Location A data displayed a cyclical pattern. Location B data displayed a trending and cyclical pattern.

Time series decomposition was then performed (Fig. 2). A decomposition of the Location A series shows no trending or seasonality, but demonstrates cyclical behavior. The absence of these structural components demonstrated the existence of noise in the time series that could worsen the performance of the forecasts.

A decomposition of the Location B series shows trending and cyclical behavior. The inclusion of these structural components could aide in the performance of the forecasts.

Understanding if this time series is white noise, exhibits random walk, or any nonlinear stochastic characteristics, is important to understanding any model's ability to forecast. A time series which is white noise will exhibit no temporal dependence structure that is learnable by the methods available. The time series is not white noise if any of the following conditions are true: [2].

- The series has a non-zero mean
- The variance changes over time
- Values correlate with lag values

Location A and B both demonstrate a non-zero mean indicating neither time series is white noise (Fig. 3). A white noise time series would shows a zero mean and standard deviation of 1.

An augmented Dickey-Fuller (ADF) Test was performed on both data sets (Fig. 3). ADF tests the null hypothesis that a unit root is present in a time series sample, and if so, the series is non-stationary. The location A and B training series show temporal dependence and are not stationary. This indicates that the process generating the series needs to be transformed prior to creating predictions.

In addition, the process generating the location B series demonstrates a stochastic or deterministic nonlinear trend and has potential structure which must be transformed before it can be learned when applying traditional approaches. Figure 4 shows two transformations on Location A and B data to account for the non-stationary and stochastic non-linear characteristics and the trending components. The decomposed time series continues to demonstrate no seasonal characteristics.

Both Location A and B's training series auto-correlation plots (Fig. 5) show values correlating with lagged values and there are spikes above the 95% confidence level. For Location A and Location B the decay is not linear indicating that the series is not a random walk and exhibits possible stochastic nonlinear behavior.

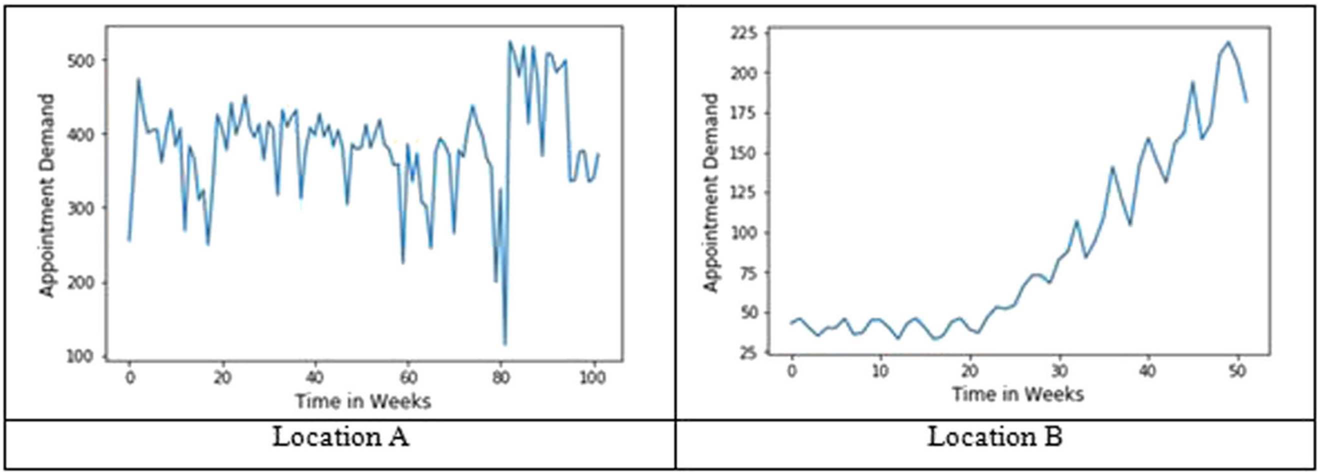

Fig. 1 Raw Time Series (Location A and Location B)

Data processing

After grouping the data into weekly cycles, a split of the data was needed to allow training to happen on a subset of the data. Data were then partitioned into training, validation, and hold-out sets, maintaining temporal sequence order. The resulting model was applied to the holdout, or out of sample, dataset to judge the effectiveness of the method (Fig. 6). This rolling window approach creates a validated method that contains the most recent information for learning signals and achieving highest accuracy on n-period ahead forecasts.

Model fitting and evaluation

A baseline to compare against was initially established. The persistence, or naïve, modeling approach simply takes the previous value in the series and uses that as the new forecast. The results are often not optimal, but establish a benchmark to improve upon. Figure 7 shows the output of the persistence model.

Model performance

With this baseline established, 19 models using various machine learning, traditional, and hybrid methods were tested against the same evaluation metric, Forecast Standard Error (FSE). This metric written in two ways is found in Eqs. 1 and 2 [4].

$$FSE = \sqrt{RMSE^2 + SE\ (Forecasted\ Y)} \tag{1}$$

$$SE(Forecast\ Error) = SD_{Errors} \cdot \sqrt{1 + \frac{1}{n} + \frac{\left(X_{value} - \overline{X}\right)}{\sqrt{n} \cdot SD^X}} \tag{2}$$

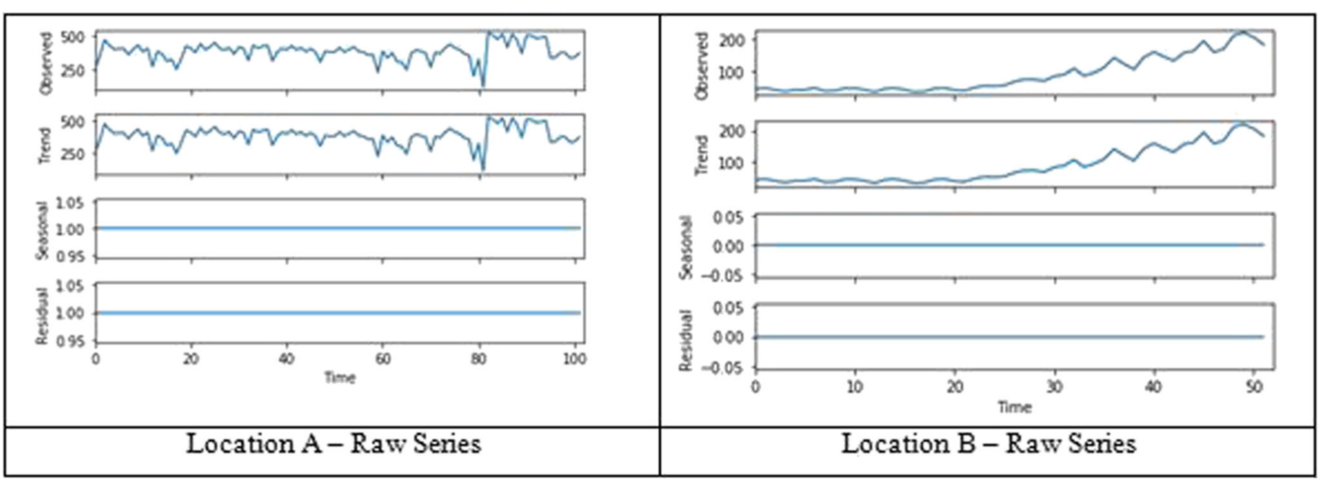

Fig. 2 Time Series Decomposition (Location A and Location B)

Location A	Location B
demand count 79.000000 mean 377.316456 std 50.566672 min 225.000000 25% 363.000000 50% 385.000000 75% 408.500000 max 474.000000	demand count 40.000000 mean 61.200000 std 30.443895 min 33.000000 25% 40.000000 50% 46.000000 75% 75.500000 max 142.000000
ADF Statistic: -2.796361 p-value: 0.058795 Critical Values: 1%: -3.525 5%: -2.903 10%: -2.589 ***series is not stationary -- is time dependent***	ADF Statistic: 1.531319 p-value: 0.997636 Critical Values: 1%: -3.639 5%: -2.951 10%: -2.614 ***series is not stationary -- is time dependent***

Fig. 3 Summary of Training Series Statistics (Location A and Location B)

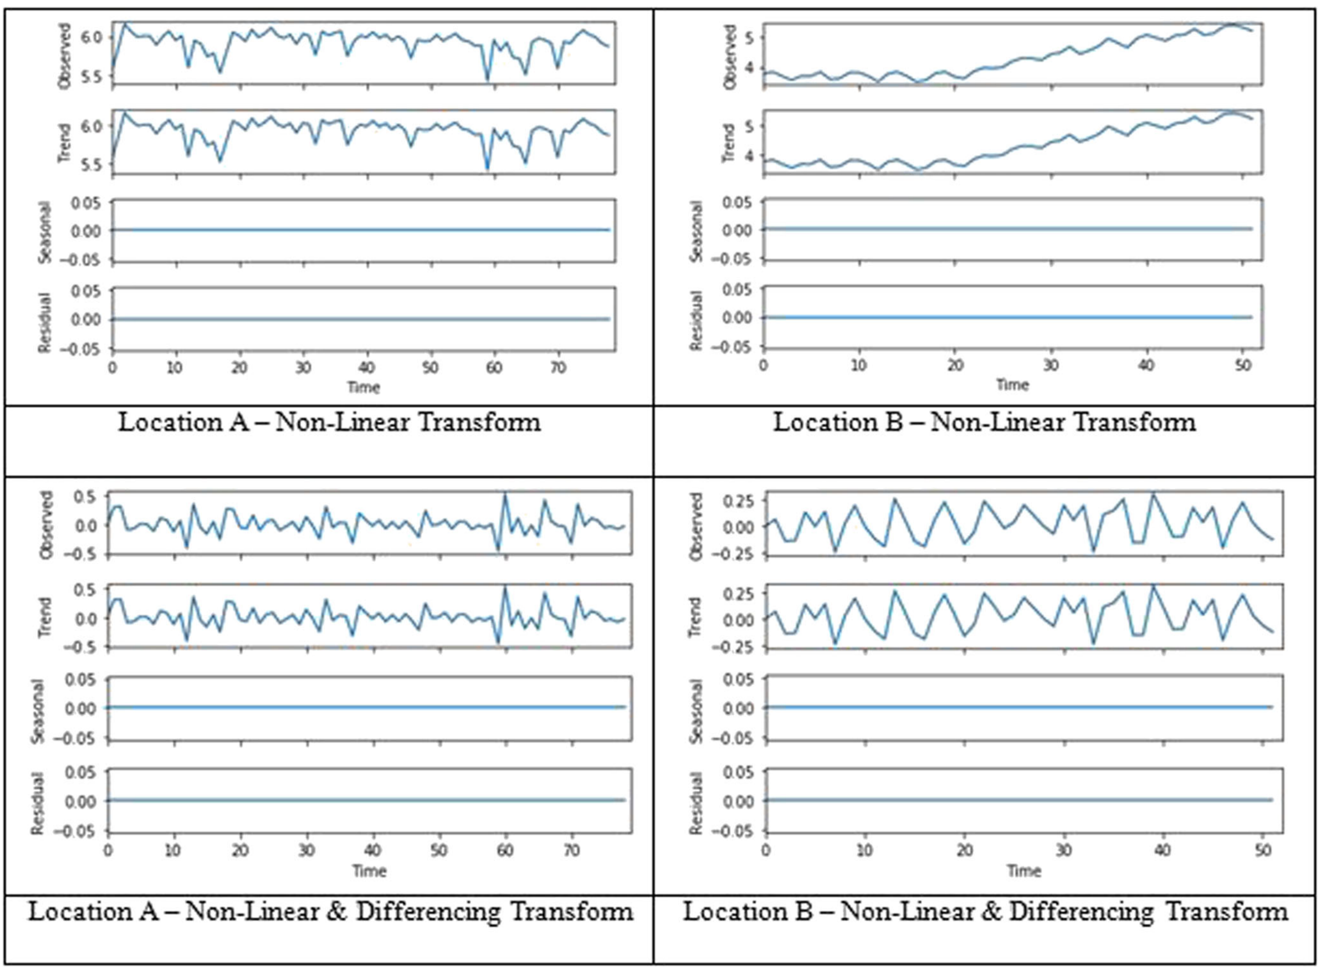

Fig. 4 Time Series Decomposition after Transformation (Location A and Location B)

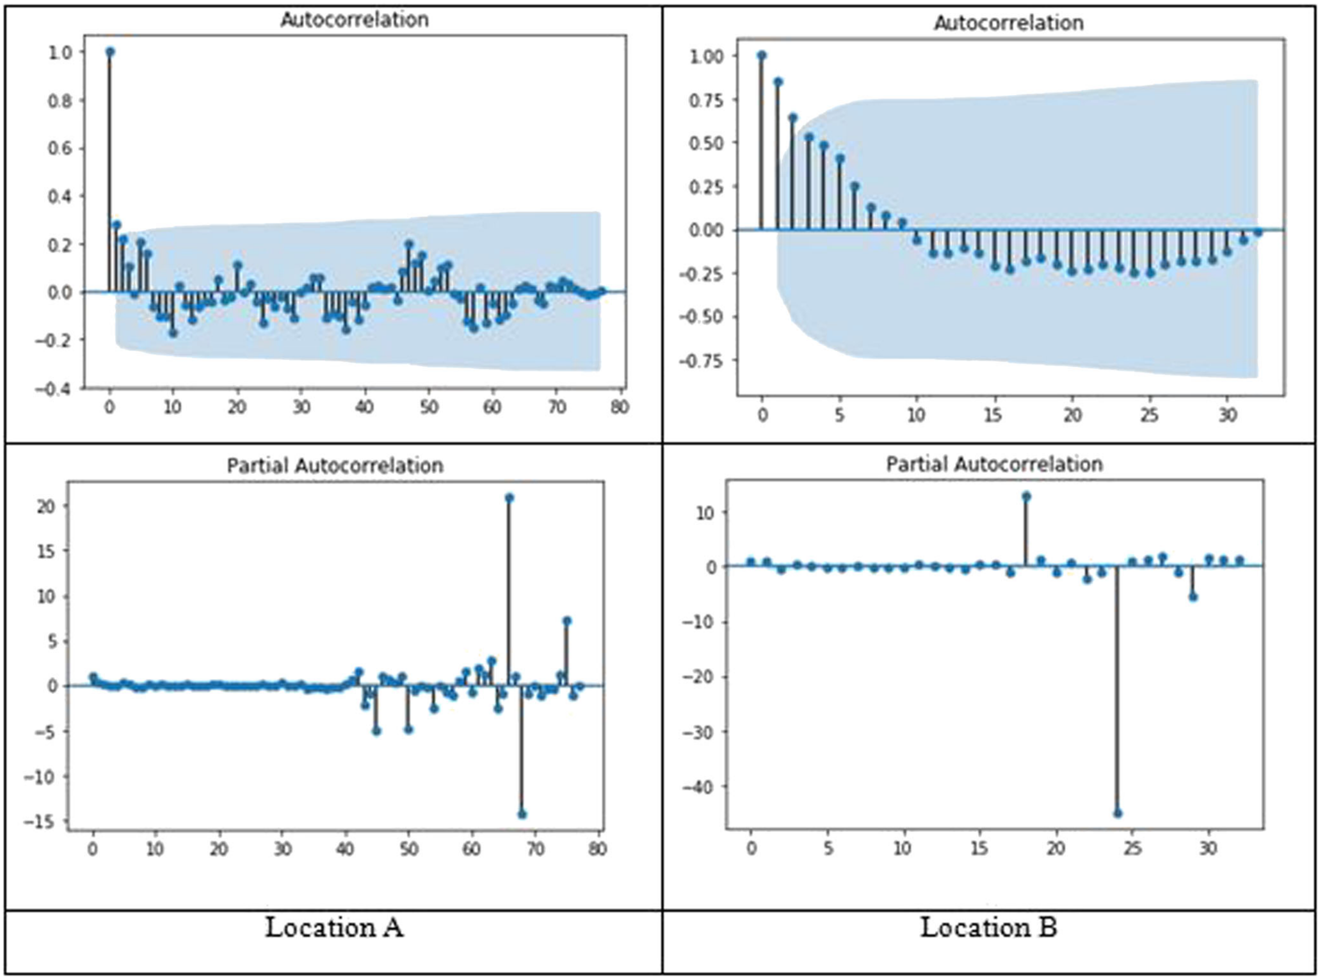

Fig. 5 Auto-Correlation Plots (Location A and Location B)

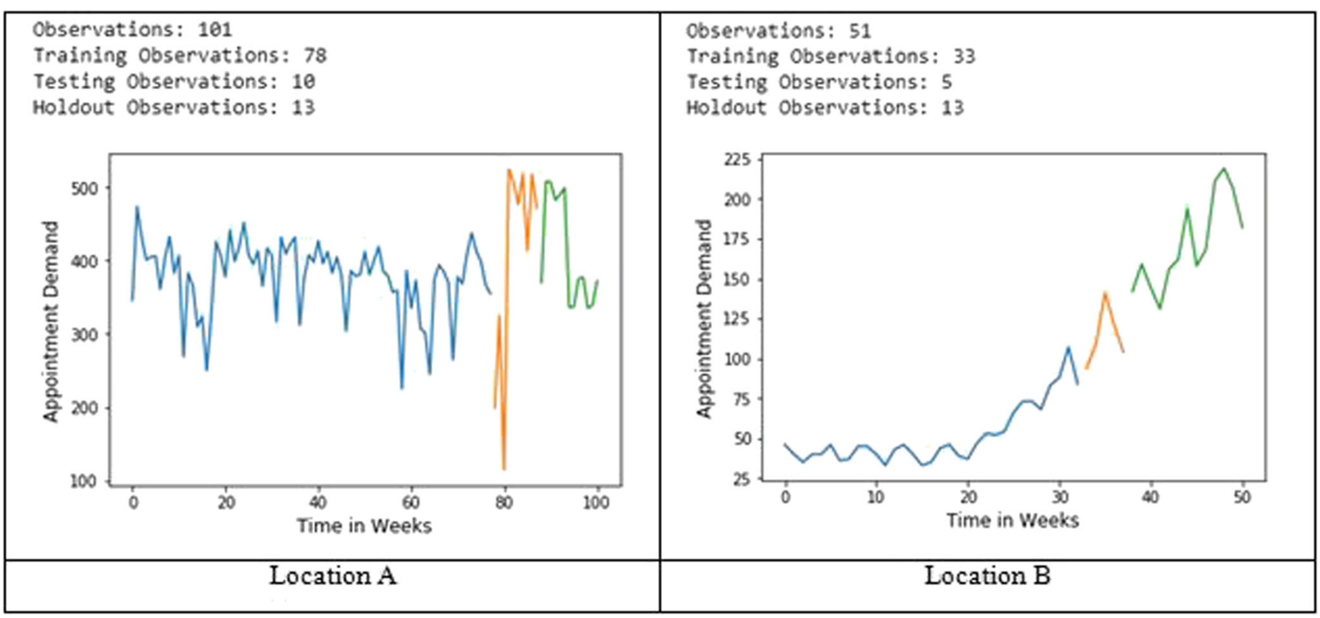

Fig. 6 Training and Validation Split (Location A and Location B)

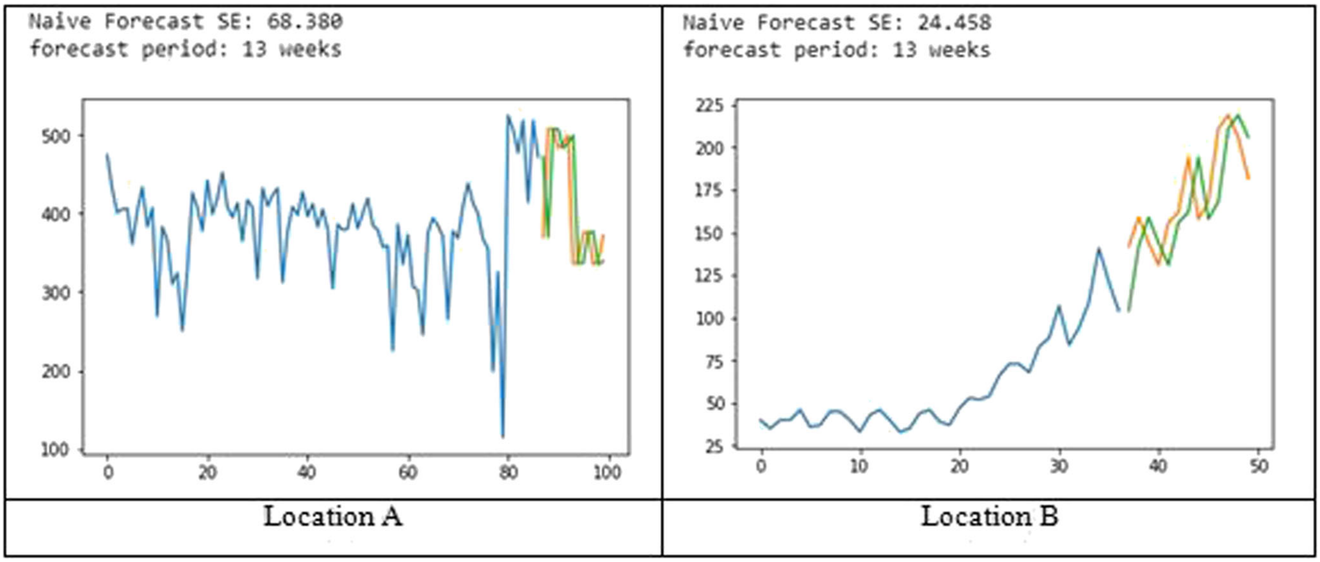

Fig. 7 Persistence (naïve) model (Location A and Location B)

Results

Location A demonstrates a cyclic, non-trending data pattern. For Location A's 39,205 patient appointment requests, models yielded a wide spectrum of results (Table 2). The naïve forecast model produced an out of sample FSE of 68.38 for a 13 week forecast horizon. The XGBoost model (machine learning) with feature engineering (additional dataset specific predictors) achieved a 94% improvement with an out of sample forecast FSE of 4.18 (vs. FSE of 5.31 without additional predictors). The ensemble model (hybrid) with feature engineering (additional predictors)

Table 2 Location A Model Performance Comparison

Model	Method	OUT OF SAMPLE FORECAST SE (FSE)
XGBoost (with additional predictors)	Machine Learning	4.18
XGBoost (without additional predictors)	Machine Learning	5.31
Ensemble (with additional predictors)	Hybrid	18.58
Ensemble (without additional predictors)	Hybrid	19.96
Linear Regression	Traditional	36.32
Grid Search CNN	Machine Learning	36.90
Grid Search LSTM	Machine Learning	47.78
Box-Cox Transformed Arima	Traditional	59.73
Grid Search Arima	Traditional	64.15
Grid Search Sarima	Traditional	65.23
Grid Search Naive and Average Forecast	Traditional	66.16
Grid Search Triple Exponential Smoothing Model (ETS)	Traditional	66.95
Naïve Model	Traditional	68.38
Moving Average	Traditional	76.67
Convolutional Neural Network (CNN)	Machine Learning	88.43
Autoregressive	Traditional	90.36
MLP Neural Network	Machine Learning	93.49
Grid Search MLP Neural Network	Machine Learning	94.51
Residual Error	Traditional	106.38
Convolution-LSTM Neural Network	Machine Learning	128.93
K-nearest Neighbor	Traditional	150.08
LSTM Neural Network	Machine Learning	192.88

Table 3 Location B Model Performance Comparison

Model	Method	OUT OF SAMPLE FORECAST SE (FSE)
XGBoost (with additional predictors)	Machine Learning	7.45
Ensemble (with additional predictors)	Hybrid	15.36
Linear Regression	Traditional	16.71
Autoregressive	Traditional	19.65
Ensemble (without additional predictors)	Hybrid	19.77
Residual Error	Traditional	22.36
Grid Search Arima	Traditional	23.00
Naïve Model	Traditional	24.46
XGBoost (without additional predictors)	Machine Learning	26.99
Box-Cox Transformed Arima	Traditional	29.33
Grid Search CNN	Machine Learning	32.30
Moving Average	Traditional	44.37
K-nearest neighbor	Machine Learning	62.90
Grid Search Sarima	Traditional	63.92
Grid Search Naive and Average Forecast	Traditional	66.16
Grid Search Triple Exponential Smoothing Model (ETS)	Traditional	66.95
Grid Search LSTM	Machine Learning	72.68
Convolutional Neural Network (CNN)	Machine Learning	85.37
MLP Neural Network	Machine Learning	92.95
Grid Search MLP Neural Network	Machine Learning	94.95
Convolution-LSTM Neural Network	Machine Learning	121.17
LSTM Neural Network	Machine Learning	190.37

followed as a second-place finisher with an FSE of 18.58 (vs. FSE of 19.96 without additional predictors), which was a 73% improvement.

Location B demonstrates a trending and cyclical data pattern. The series was not white noise. For Location B's 4,538 patient appointment requests, models yielded a narrower spectrum of results (Table 3). The naïve forecast model produced an out of sample FSE of 24.46 for a 13-week forecast horizon. The XGBoost model with feature engineering (additional dataset specific predictors) achieved a 70% improvement with an out of sample forecast FSE of 7.45 (vs. FSE of 26.99 without additional predictors). The Ensemble model with feature engineering (additional predictors) achieved a 37% improvement with an out of sample forecast FSE of 15.36 (vs. FSE of 19.77 without additional predictors).

Discussion

Machine learning approaches are often thought of as the next evolution in data science. However, when comparing 20 machine learning, hybrid, and traditional models on these two substantially different datasets, it was determined that advanced machine learning approaches alone did not always deliver the highest levels of accuracy. In Location A with cyclical non-trending data, of the top half of performing models, only 36% involved pure machine learning, but the optimum model was indisputably a machine learning method. In Location B with trending and cyclical data, of the top half of performing models, only 27% involved pure machine learning and although the results were more tightly clustered, the computationally intensive and complicated methods of machine learning often failed to outperform a simpler traditional method.

The differences in results between locations A and B highlight the importance of careful data set analysis prior to choosing a forecasting model. For instance, traditional methods can outperform machine learning with linear data structures whereas innovative feature engineering can markedly improve a machine learning method's ability to locate the right predictors for generating enhanced results.

XGBoost (machine learning) is a powerful highly scalable algorithm, which drives fast learning through parallel and distributed computing and offers efficient memory usage. In contrast to bagging techniques like Random Forest, in which trees are grown to their maximum extent, boosting makes use of trees with fewer splits. The smaller trees produced by XGBoost (machine learning) are not very deep and therefore highly interpretable. XGBoost benefits from feature engineering to transform the inputs of the raw time series into predictable components. Even without application of feature engineering with additional dataset specific predictors, XGBoost works well on series that exhibit cyclical patterns (Location A), but as expected performs less optimally on a more complex cyclical and trending dataset (Location B) where it underperformed the naïve method. Feature engineering with the inclusion of additional predictors selected based on the dataset (in this case: Day of the month, Day of the week, Ordinal date, and Hour of the day) significantly enhances the forecasting value of the model, particularly on Location B's more complex dataset, transforming it into the best performing algorithm.

Ultimately to provide a consistently good solution across multiple datasets, without the need for additional dataset specific predictors, great promise lay in the use of hybrid methods which combine the strengths a traditional method with a machine learning method. The ensemble model used in this study combines a machine learning method (XGBoost) which is good at cyclical behavior with a traditional method (linear regression) that handles trending well. The results of this hybrid method on both Location A and B's datasets, without additional dataset specific predictors, were consistent clustered near the best performers and demonstrated even further improvement after adding feature engineering with additional predictors to the machine learning component, indicating the usefulness and flexibility of hybrid methods in a variety of datasets with different characteristics.

Conclusion

In order to obtain the most accurate forecasting and optimal solution, data characteristics should be first examined. Various traditional, hybrid, and machine learning methods should be tested and compared for the best result. Use of intelligent feature engineering with machine learning methods should be considered as it can dramatically improve results. Finally, hybrid methods perform well on many different types of datasets by utilizing the strengths of both traditional and machine learning methods.

Compliance with ethical standards

Conflict of interest The authors declare that they have no conflict of interest.

Clinical relevance statement This study aims to accurately forecast patient appointments by comparing various machine learning, traditional and hybrid methods. The step-by-step approach discussed can help clinics to develop their own patient appointment predictive model, and thus effectively plan the appropriate resources and improve patient satisfaction.

Protection of human and animal subjects Only retrospectively collected data were used, and no patients were contacted for the study. The study did not involve any protected health information (PHI) and all data points were de-identified.

Informed consent Informed consent was obtained from all individual participants included in the study.

References

1. DeLurgio, S., Denton, B., Cabanela, R. et al., Forecasting weekly outpatient demands at clinics within a large medical center. Production and Inventory Management Journal 45(2):35–46, 2009.
2. Brownlee, J., Deep Learning for Time Series Forecasting, 2018. Retrieved from https://machinelearningmastery.com/deep-learning-for-time-series-forecasting/
3. Jain, A., What I learnt about Time Series Analysis in 3 hour Mini DataHack?, 2016. https://www.analyticsvidhya.com/blog/2016/02/hand-learn-time-series-3-hours-mini-datahack/
4. How to Compute the Estimated SEs of Forecasted Y and the Forecast Error. www3.wabash.edu/econometrics/econometricsbook/.../HowToFindSEForecast.doc
5. XGBoost developers. Introduction to Boosted Trees. XGBoost, 2016. https://xgboost.readthedocs.io/en/latest/tutorials/model.html. Accessed 28 December 2018
6. Abdus-Salaam, H., Davis, L., de Oliveira Mota, D., Modeling dependent demand arrivals within an open-access scheduling system. Paper presented at: Science and Technology for Humanity (TIC-STH), 2009 IEEE Toronto International Conference, 2009.
7. Abraham, G., Byrnes, G. B., and Bain, C. A., Short-term forecasting of emergency inpatient flow. IEEE transactions on information technology in biomedicine: a publication of the IEEE Engineering in Medicine and Biology Society. 13(3):380–388, 2009.
8. Afilal, M., Yalaoui, F., Dugardin, F., Amodeo, L., Laplanche, D., and Blua, P., Forecasting the Emergency Department Patients Flow. J Med Syst. 40(7):175, 2016.
9. Ahmadi-Javid, A., Jalali, Z., and Klassen, K. J., Outpatient appointment systems in healthcare: A review of optimization studies. European Journal of Operational Research. 258(1):3–34, 2017.
10. Akyuz, A. O., Uysal, M., Bulbul, B. A., Uysal, M. O., Ensemble approach for time series analysis in demand forecasting: Ensemble learning. Paper presented at: 2017 IEEE International Conference on Innovations in Intelligent SysTems and Applications (INISTA), 2017.
11. Al-Haque, S., Ceyhan, M. E., Chan, S. H., and Nightingale, D. J., Responding to traveling patients' seasonal demand for health care services. Mil Med. 180(1):111–117, 2015.
12. Bouckaert, N., Van den Heede, K., and Van de Voorde, C., Improving the forecasting of hospital services: A comparison between projections and actual utilization of hospital services. Health Policy. 122(7):728–736, 2018.
13. Boutsioli, Z., Measuring Unexpected Hospital Demand: The Application of a Univariate Model to Public Hospitals in Greece. Hospital Topics. 87(4):14–21, 2009.
14. Calegari, R., Fogliatto, F. S., Lucini, F. R., Neyeloff, J., Kuchenbecker, R. S., and Schaan, B. D., Forecasting Daily Volume and Acuity of Patients in the Emergency Department. Computational and Mathematical Methods in Medicine. 2016: 3863268, 2016.
15. Cheng, C. Y., Chiang, K. L., and Chen, M. Y., Intermittent Demand Forecasting in a Tertiary Pediatric Intensive Care Unit. J Med Syst. 40(10):217, 2016.

16. Claudio, D., Miller, A., and Huggins, A., Time series forecasting in an outpatient cancer clinic using common-day clustering. IIE Transactions on Healthcare Systems Engineering. 4(1):16–26, 2014.
17. Côté, M. J., Smith, M. A., Eitel, D. R., and Akçali, E., Forecasting Emergency Department Arrivals: A Tutorial for Emergency Department Directors. Hospital Topics. 91(1):9–19, 2013.
18. Côté, M. J., and Tucker, S. L., Four methodologies to improve healthcare demand forecasting. Healthcare Financial Management. 55(5):54–58, 2001.
19. Heroman, W. M., Davis, C. B., Farmer Jr, K. L., Demand forecasting and capacity management in primary care. Physician executive. 38(1):30–34, 2012. http://ovidsp.ovid.com/ovidweb.cgi?T=JS&PAGE=reference&D=emed13&NEWS=N&AN=369644991
20. Hong, W. S., Haimovich, A. D., and Taylor, R. A., Predicting hospital admission at emergency department triage using machine learning. PloS one. 13(7):e0201016, 2018.
21. Kadri, F., Harrou, F., Chaabane, S., and Tahon, C., Time Series Modelling and Forecasting of Emergency Department Overcrowding. J Med Syst. 38(9):1–20, 2014.
22. Luo, L., Luo, L., Zhang, X., and He, X., Hospital daily outpatient visits forecasting using a combinatorial model based on ARIMA and SES models. BMC Health Serv Res. 17(1):469, 2017.
23. Mai, Q., Aboagye, S. P., Sanfilippo, F. M., Preen, D. B., and Fatovich, D. M., Predicting the number of emergency department presentations in Western Australia: A population-based time series analysis. Emergency Medicine Australasia. 27(1):16–21, 2015. https://doi.org/10.1111/1742-6723.12344.
24. Ouyang, Y., and Yin, H., Multi-Step Time Series Forecasting with an Ensemble of Varied Length Mixture Models. Int J Neural Syst. 28(4):1750053, 2018.
25. Schweigler, L. M., Desmond, J. S., McCarthy, M. L., Bukowski, K. J., Ionides, E. L., and Younger, J. G., Forecasting models of emergency department crowding. Acad Emerg Med. 16(4):301–308, 2009.
26. Shan, J., Zhang, H., Liu, W., and Liu, Q., Online Active Learning Ensemble Framework for Drifted Data Streams. IEEE Transactions on Neural Networks and Learning Systems.:1–13, 2018. https://doi.org/10.1109/TNNLS.2018.2844332.
27. Sun, Y., Heng, B. H., Seow, Y. T., and Seow, E., Forecasting daily attendances at an emergency department to aid resource planning. BMC Emergency Medicine. 9:1, 2009.
28. Villani, M., Earnest, A., Nanayakkara, N., Smith, K., de Courten, B., and Zoungas, S., Time series modelling to forecast prehospital EMS demand for diabetic emergencies. BMC Health Services Research. 17:1–9, 2017.
29. Yu, L., Hang, G., Tang, L., Zhao, Y., and Lai, K., Forecasting Patient Visits to Hospitals using a WD&ANN-based Decomposition and Ensemble Model. Eurasia Journal of Mathematics, Science and Technology Education. 13(12):7615–7627, 2017.