Forecasting Crew Member Scheduling for Jakarta Emergency Ambulance Service (118)

Abhishek Kumar
S.No: 23039854

Introduction And Problem Definition

Jakarta ambulance services operate 24 hours a day, seven days a week.
Call data is provided from March 1, 2019 to September 30, 2019. call data in the format of date and time.
We will inspect different forecasting techniques, namely Naïve baseline, mean, SES, simple linear regression, Holt linear, Holt's winter, and the and the Arima model, to find the best method for the next 2 months of forecasting. So that maximum utilisation of resources and services, accordingly schedule a roaster for their crew member.

Data Cleaning

Inspection of dupicate and missing values.

Call data counts	After removing Duplicate value	Missing values	Final counts
32905	32709	0	32709

Numerical Summaries and Graphical Summaries

statistics	Per days call data	24 hr call data
Resample_data_counts	214.00	5136.00
minimum	60.00	0.00
maximum	210.00	26.00
mean	152.85	6.37
25%	119.00	2.00
Median (50%)	163.00	5.00
75%	183.0	10.00
Standard Deviation	36.10	4.97
Variance	1368.56	24.73
Mean Absolute Deviation	28.00	3.00
Mean of Squared Deviations	1362.17	24.73

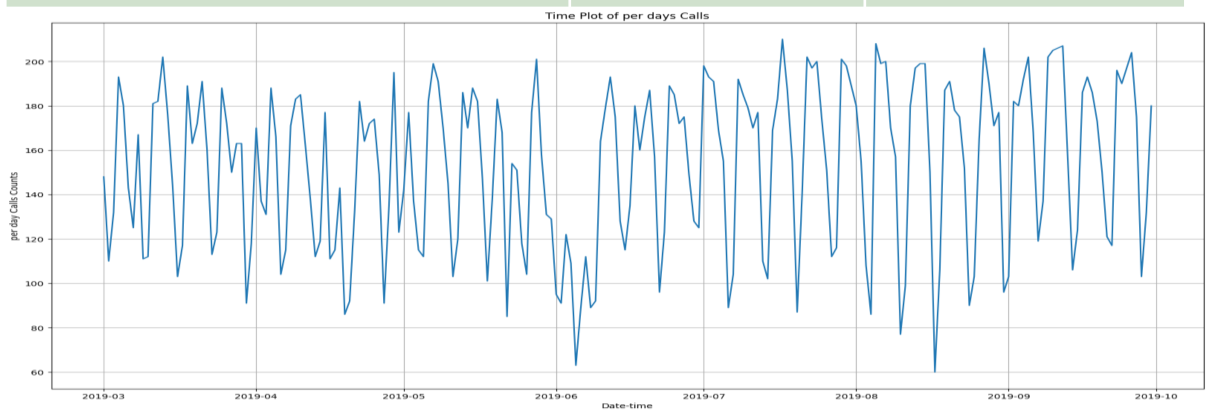

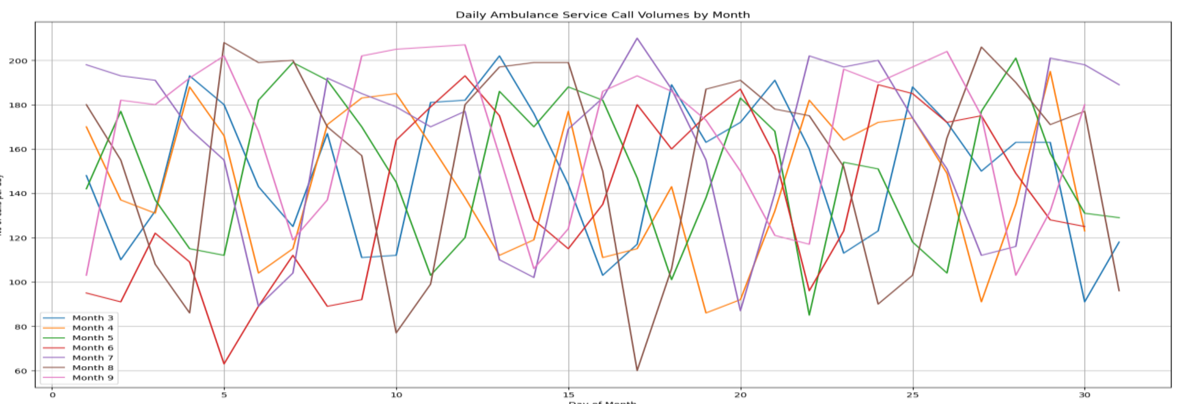

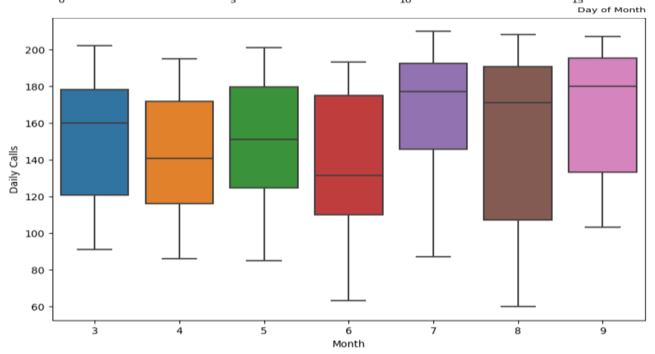

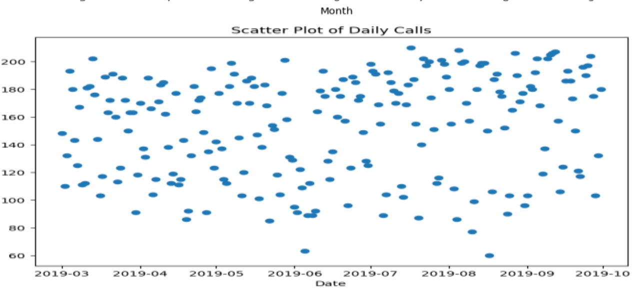

Data is very volatile and suddenly changes from one point to another. We sample data on a 24-hour and daily basis, and we can see comparisons between monthly calls and hourly calls. I identifies the busy time for calls is from 9 a.m. to 3 p.m. There is no missing value.
We will use per-day call resampling data to find the best model to forecast two months ahead. If we look at the at the time plot for whole data from March 2019 to September 2019, no trend or pattern can be found directly. Also, there is no trend or seasonality if we see a monthly comparison for 30 days.

Data Exploration

Variance is much fluctuating over time, with a mean also sudden fluctuation between June and July and August and September. To achieve stability in the data, I decided to apply a log transformation. We can see in the graph that variance seems to be approximately stable with means. There is very little fluctuation in variance, so the data has low variability.

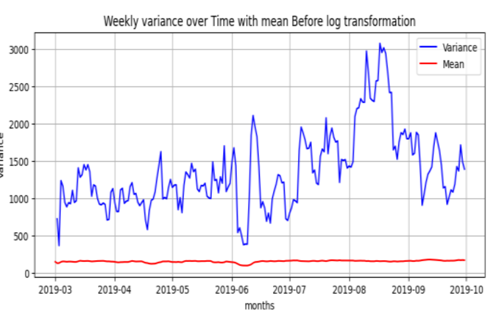

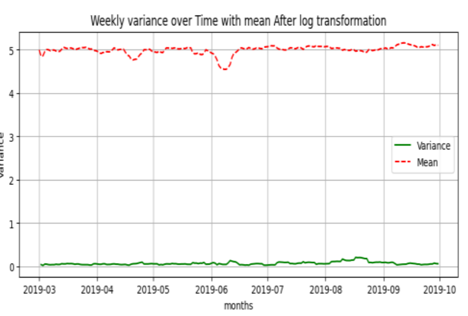

Decomposition

To identity Trends and seasonality in the data we inspected by decomposition plot. There is no clear evidence showing the above variability of data because residuals are still fluctuating. This decomposition plot suggests that the time series data contains both seasonal patterns and random fluctuations, with no clear long-term trend. Seasonal patterns look additive because there is no exponential growth in the data shown. Further analysis may be needed to identify the specific factors driving the variability in the data.

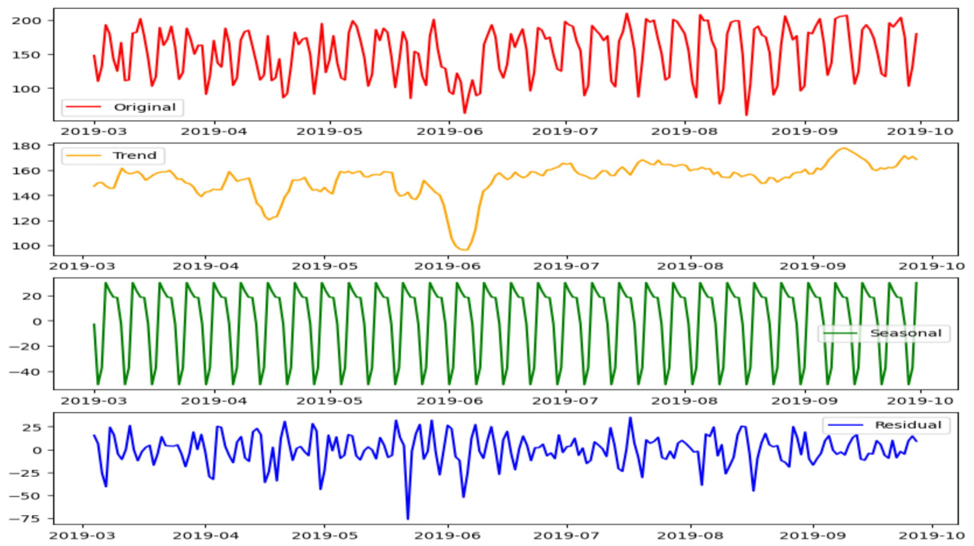

Moving Average

Data is recorded on a daily basis, so we will use 7MA, double MA, and weighted MA to find any trend that exists, either an increase or decrease.
7MA is smoothed as compared to actual data, and Double 7MA is more smoothed than 7MA. By using these two moving averages, no trend was found. So apply weighted MA to inspect further. Additionally, try to check 30MA to identify trends on a monthly basis since we have 7-month data. The moving average graph does not show clear increase or decrease trends. It looks to increase at one point and decrease at another. Overall, no trend was found.

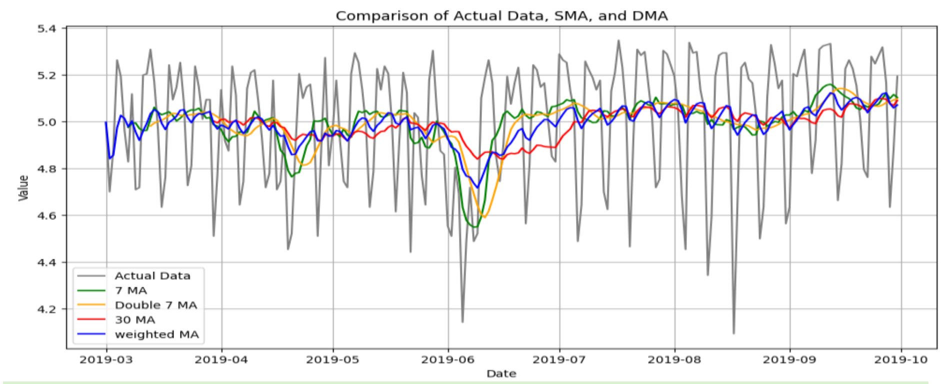

STL Decomposition

The STL decomposition graph also shows seasonality but does not show any trend. Now we can confidently say that there are no trends in the data.

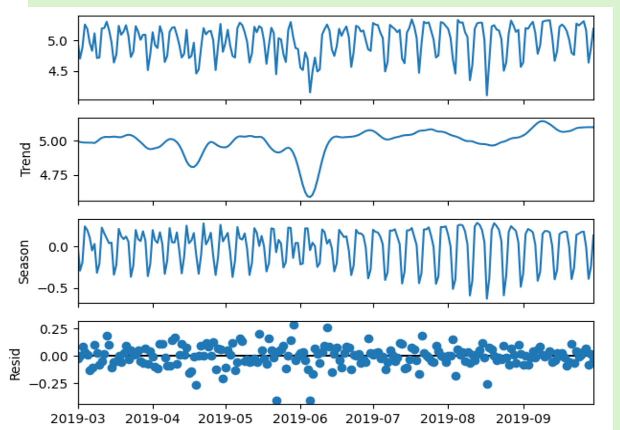

STL is able to handle data where there is no quarterly or monthly seasonality; due to this, we got weekly or three-month seasonality in data.
Residual is depicting untraced and random fluctuation of per day calls data.

Baseline, SES, holt Linear, holt's winter, regression Models

We split the data between training and testing the model, like the first 5 months of data for training and the last 2 months of testing, which is approximately 70% and 30%, respectively.
Training data: 01/03/2019 to 31/07/2019
Testing data: 01/08/2019 to 30/09/2019
We implemented the following models:
The Naïve method is used as a baseline model for comparison to others. Along with Naïve, we implement the mean model as a baseline model. Other methods include SES, Holt linear, Holt's winter additive method, and linear regression. Out of these, Holt's winter additive method captures seasonality well as compared to other methods. The reason behind this method is that it works well with both trend and seasonality. In our call data, seasonality exists, which was confirmed by additive decomposition and the STL decomposition method.
All models train over log-transformed values, and forecasted values are reversed to the original scale. So the comparison of models in the graph is in original scale values. We can see how it behaves with the actual test data in the graph below.

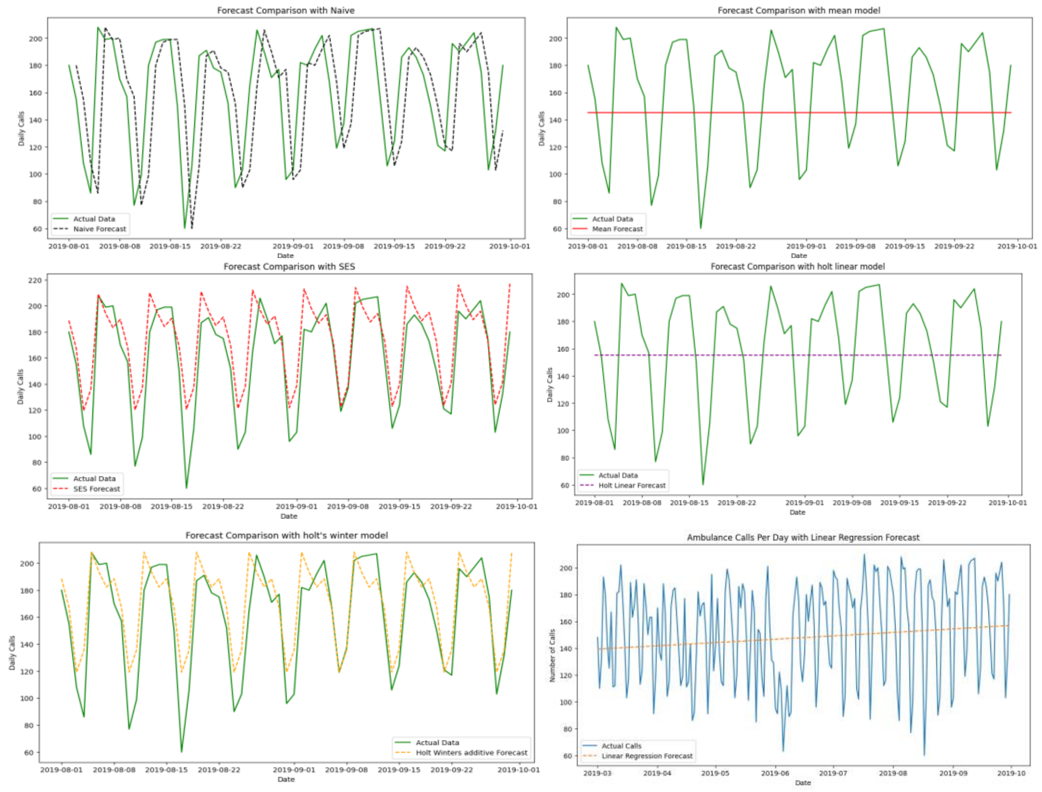

ARIMA Model

We applied the Dicky Fuller test to call-transformed data to check if the data was stationary or not. I found that the data was non-stationary. Here is the result of the Dicky Fuller test. After applying the first order of differencing, I got stationary data. When I see the above test result, the ADF statistic of transformed data is greater than all critical values, and the p-values are greater than zero. On the other hand, ADF statistics for 1st order differencing are less than all critical values, and the p-value is too low to be zero, which shows 1st order differencing needs to be implemented in the Arima model.

Critical values

1%	5%	10%
-3.464	-2.876	-2.575

Dicky fuller test Result

	Transformed data	First order differencing data
ADF	-2.364	-6.202
P-value	0.152	5.766 x (10)^-8

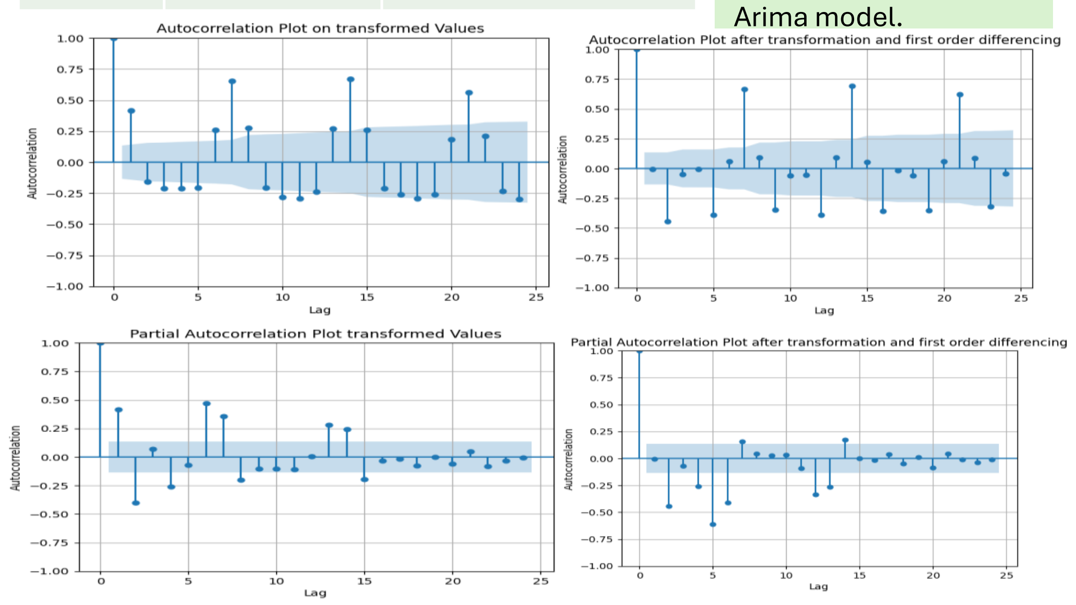

ACF plots significant autocorrelation at several lags, which we notice at Lag1 and Lag14. In PACF plots, there is a significant spike at Lag 1, which shows a rapid decrease in partial autocorrelation. Plots indicate the AR(1) model should be used. However, after inspection by fitting, AR(14) is given a very low MSE as compared to AR(1). The reason is that the significant positive autocorrelation at lag 14 in the ACF plot suggests the presence of a seasonal pattern in the data, which also came from the decomposition method. An AR(14) model captures this seasonal behaviour more explicitly, which leads to better model performance. We fitted ARIMA(1,1,0), ARIMA(5,1,0), ARIMA(12,1,0), and (14,1,0) with AIC: (43.61, -44.02, -81.52, -99.01) and BIC: (49.66, -25.87, -42.21, -53.65) respectively. The ARIMA model with order (14, 1, 0) has the lowest AIC and BIC values, indicating that it provides the best balance between goodness of fit and model complexity. We chose ARIMA (14, 1, 0) for farecasting, which has a lower minimum MSE than other models.
After implementing the ARIMA model, I found the Arima model for forecasting more accurate as compared to different implemented models. There is very low MSE and MAPE for ARIMA as compared to others. ARIMA more accurately captures patterns and even more volatile data. A comparison of all implemented models is shown below:

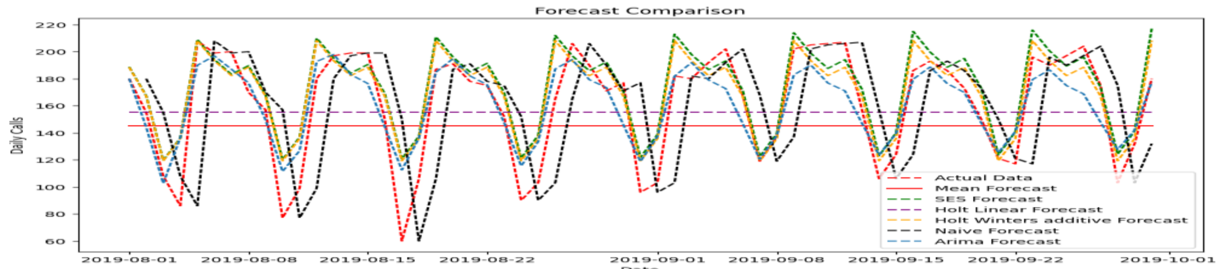

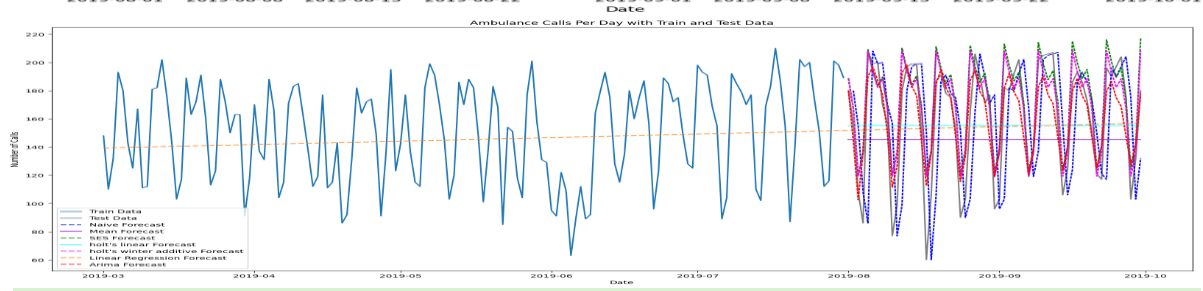

Error statistics of all Models

	Transformed forecasted data		Transformed forecasted data reversed to original Scale	
models	MSE	MAPE	MSE	MAPE
Naive	0.106521	0.047350	1847.766667	0.240027
Mean	0.090513	0.053746	1852.377618	0.240027
SES	0.030812	0.025584	480.546753	0.138997
Holt linear	0.086573	0.049565	1641.839874	0.266979
Holt winter	0.027788	0.023836	414.887429	0.128294
Linear regression	0.070597	0.045912	1351.140048	0.239554
ARIMA	0.026031	0.023134	395.877781	0.120379

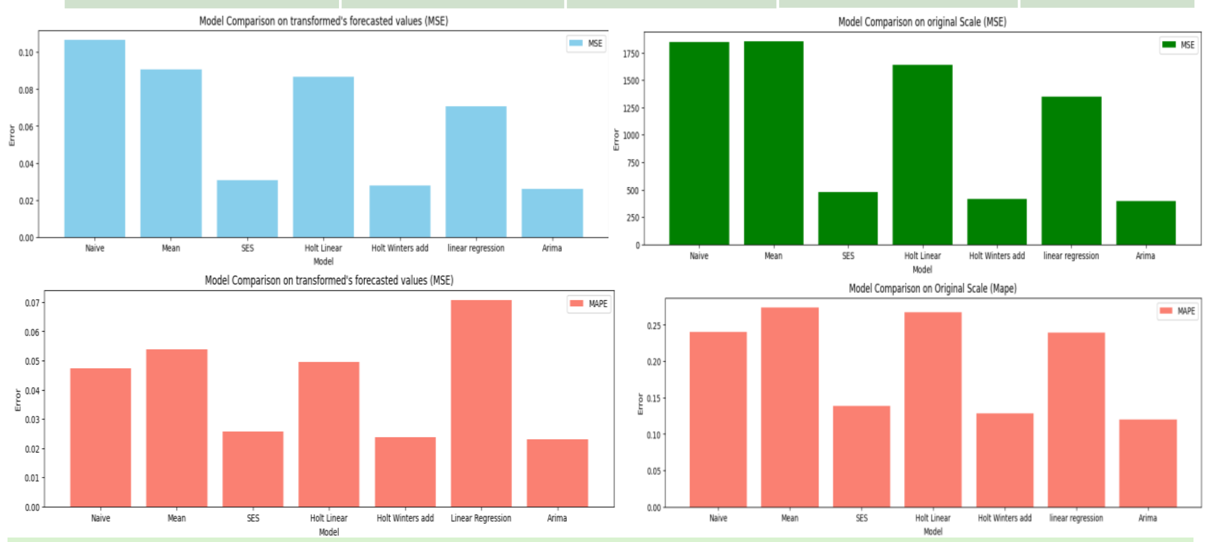

forecast Summary

Using Arima models per days call forecasted values (reversed to original Scale) from 01/10/2019 to 30/11/2019 : 185.43, 174.24, 167.58, 148.02, 129.56, 143.41, 176.55, 183.52, 173.21, 166.80, 148.52, 131.43, 144.34, 175.37, 182.41, 172.40, 165.78, 148.85, 133.22, 145.54, 174.34, 180.86, 171.64, 165.13, 149.25, 134.82, 146.45, 173.45, 179.72, 170.84, 164.34, 149.64, 136.42, 147.41, 172.51, 178.47, 170.20, 163.76, 149.99, 137.84, 148.28, 171.75, 177.36, 169.48, 163.17, 150.39, 139.24, 149.08, 170.95, 176.30, 168.89, 162.66, 150.72, 140.52, 149.86, 170.26, 175.29, 168.29, 162.21, 151.09, 141.73

Conclusion

The Arima model is the best-fitted model to capture past data patterns closely with very low MSE and MAPE. Holt's winter additive model is the second-best-fitted model, and MSE and MAPE are much closer to the ARIMA model because this method is able to handle seasonality. There is a lot of variability in the data. The variance was not stable. After inspecting with different methods, we did not trace any trend but seasonality, which was not on a quarterly or monthly basis. We explored data on the basis of per-day call counts, but it can be explored more to gain some insight into hourly calls with the use of different methods.