

When does Machine Learning work in Time Series Forecasting?

Sukhdeep Singh, Gaurangi Gupta

Abstract— The study mainly focuses on the applications of Machine Learning techniques like Artificial Neural Network and Random Forests in time series forecasting which is a sequence of observations measured over time. We have simulated a time series and obtained predictions using Machine Learning methods like ANN and Random forest algorithms to check the accuracy and robustness of machine learning methods in making the predictions. Here we also compare the performance of these two methods.

Index Terms— accuracy, ANN, Machine learning, Random forest, Time series

1 INTRODUCTION

THIS paper involves simulation of a time series which will make predictions using Machine Learning algorithms where we have used ANN and Random Forest algorithms.

1.1 Objective of the Study

We want to see how well these methods work in autoregressive time series with additional stationarity. We will test these methods on simulated data which is generated according to a true model. This true model has a trend component, seasonality component and the error term. We have used monthly data as monthly time series are the most common type of data in businesses and economics and they are also difficult to predict as compared to other seasonal time series such as the quarterly time series. Machine learning methods use information and patterns regarding historical values to predict future activity.

1.2 Significance of Study

The classical decomposition method decomposes a time series into seasonal, trend, cyclical and irregular components. It is due to the belief that seasonality can hinder the measurement of other components and therefore need to be separated. Sometimes, the seasonal component is not independent of the non-seasonal component and therefore it is not advisable to separate the two components.

These models are parametric in nature which means that these models have some finite number of parameters associated with them. Parameters are quantities that summarize data and provide some information about the data. A limitation of these models is that the model form must be specified [qian].

A linear relationship is assumed between the values and their respective historic values. Assuming linear relationship limits the classical method's ability to model complex non-linear problems which are encountered in reality. Among the traditional predicting methods, for example the Box-Jenkins ARIMA makes use of information underlying in past values of the variable itself and errors in the past. Machine learning methods make use of artificial algorithms to learn pattern, under-

lying given data without knowledge of any kind of underlying relationships.

2 LITERATURE REVIEW

2.1 Time Series

A time series is a sequence of observations measured over time. Time series analysis is the use of a model to predict future values based on previously observed values. The objectives of time series analysis include forecasting future values. In a time series it is assumed that each point is independent of each other. A stationary time series is one which represents such a condition. In a stationary time series, properties like mean, variance and autocorrelation are constant over time.

A time series consists of four components; *trend*; which is the linear or non-linear component that changes over time and doesn't repeat itself, *seasonality*; patterns of change that occur at specific regular intervals, *cyclic fluctuations*; which correspond to periodic fluctuations in the data which is not seasonal and *irregular variations*; which are the non-random variations of the series. It is due to the belief that seasonality can hinder the measurement of other components and therefore need to be separated. Sometimes, the seasonal component is not independent of the non-seasonal component and therefore it is not advisable to separate the two components.

The multiplicative model of the time series through which we will be simulating data can be written as:

$$Y_t = T_t \cdot SI_t + E_t;$$

where T_t represents the *trend component*, SI_t represents the *seasonal variation* and E_t represents the irregular or the residual variations.

If the trend is linear, then it can be written by:

$$T_t = a + bt :$$

where ' a ' and ' b ' are constants and ' t ' is a given time.

A lag in a time series represents a fixed amount of passing time. The ' k^{th} ' lag is the time period that happened ' K ' time points before time ' i '.

2.2 Machine Learning Techniques

Machine learning is the practice of using algorithms to analyze data, learn from it and then make a prediction about something.

2.3 Artificial Neural Network

Artificial neural networks are complex multivariable statistical models for the estimation of an unknown expectation on function of a random variable ' y ' given explanatory variables ' x '. They have been proved to model any type of relationship with high accuracy.

ANN can be written as nonlinear regression models containing input variables x (which can be random variables) and output or dependent variables ' y '. These models are either used to determine unknown relations between the 2 variables.

A general type of the neural network model is the feed forward network, where the information only passes from the left to the right. It comes into the network through the input layer, goes through the hidden layers and comes out through the output layer. It does not loop between the hidden layers; neither does it loop within a single hidden layer nor in a neuron itself.

The x variables or the parameters that determine the y variable form the input layer. The output layer consists of the y variable. The input variables are weighted up and the output variable is calculated. All of the neurons from the input variables have synopsis connecting them to the hidden layer, and these synopses have weights. Some of these weights will have zero weights and some of them will have non-zero weights indicating that some weights are important whereas some of them are not that important. The output variable can be calculated as the weighted sum of all the inputs.

In time series forecasting, the input nodes are formed by the previous lagged observations, while the output layer provides the forecast for the future value. The hidden layers process the information received by the input layer with appropriate non-linear transfer function.

2.4 Random Forests

Random forests are one of the most popular and powerful machine learning algorithms used both in regression and classification. Random forest is a collection of different decision trees. The aim to random forest is to combine a set of decision trees, each of which is constructed using a bootstrap sample.

This algorithm creates the forests with a number of decision trees. More the trees in the forest, more robust is the prediction. In this technique, multiple trees are grown. To classify each object based on attributes, each tree gives its prediction. The forest chooses the classification with the most votes and in the case of regression takes the output of the different trees.

This process is known as *bagging* and it improves the stability and accuracy of the model, reduce the variance and helps avoiding overfitting.

Studies have shown that random forest does not overfit as more trees are added but produce a limiting value of the generalization error. The random forest generalization error is estimated by an out-of-bag error, i.e. the error of training points not contained in the bootstrap training sets. The algorithm of random forest for regression is [1]:

1. For $k = 1$ to K :
 - 1.1. Draw a bootstrap sample L of size N from the training data.
 - 1.2. Grow a random-forest tree T_k to the bootstrapped data, by recursively repeating the following steps for each node of the tree, until the minimum node, size m is reached.
 - 1.2.1. Select F variables at random from the n variables.
 - 1.2.2. Pick the best variable/split-point among the F .
 - 1.2.3. Split the node into two daughter nodes.
2. Output the ensemble of trees $\{T_k\}_{k=1,2,\dots,K}$.

To make a prediction at a new point x :

The two main parameters of random forest are the number of trees in the forest K and the number of input variables at each split F .

2.5 Indicator of efficiency

We used the root mean square error (RMSE) as the indicator of accuracy of the prediction. It is the square root of the mean of the difference between the actual y_t values and the values predicted by the machine learning methods. The higher the error, the less accurate is the technique.

2.6 Training, Validation and Test dataset

The dataset is generally split into the training, validation and the test set. The training set is the set through which the machine learning technique learns about the data. The algorithms find relations between variables using this data which further can be used to predict values when we have a different dataset. Different models are created with different parameter values. These parameter values are assessed by the validation set.

This is the next portion of the split. It is used to assess how well the model has been trained. This set is used to pick the best model among all the models that are trained. Using this dataset, we choose the model with the best parameters. This is indicated by different indicators like the Root Mean Square error (RMSE), Mean Square Error (MSE), Mean Absolute percentage error (MAPE) etc.

The test set tells about the accuracy of the chosen model. The training dataset cannot be used to measure the performance of

the model as the model may have memorized some of the training data and since the validation set was used to pick the best model, it may cause some bias while measuring the performance of the model. The test set behaves as an out-of-sample dataset which is completely new to the model.

2.7 Applications of Machine Learning

- Image Recognition

Image recognition is one of the most popular applications of machine learning. It is used to identify objects, persons, places etc. The most popular one of them is 'automatic friend tagging suggestion'. This is when social media applications ask the user automatically suggests the name of the person in the photograph uploaded.

- Speech Recognition

This is the process of converting voice instructions to text. Google assistant, Siri, Alexa are all using speech recognition technology.

- Product Recommendations

It is through machine learning algorithms that we get suggestions for entertainment series or movies on Netflix after we have watched a film or suggestions on any product we should buy when we search something on Amazon.

3 METHODOLOGY

3.1 Generating Time Series

A time series was simulated in order to carry out all the tests. We have simulated a monthly time series as it is very common for forecasting. Moreover, these are difficult to deal with than other types of seasonal data such as yearly or quarterly. Simulated time series was generated according to the multiplicative model:

$$Y_t = T_t \cdot SI_t + E_t,$$

where,

Y_t is a function of time ' t ' and is an observation at a given time ' t '.

$T_t = 100 + 0.6t$; is the linear trend,

SI_t is the seasonal index as given in Table 1,

E_t is the error term following a normal distribution, $N(0, \sigma^2)$.

The errors are drawn from a standard normal distribution.

A total of 228 points were generated.

Month	January	February	March	April	May	June	July	August	September	October	November	December
SI	0.75	0.80	0.82	0.90	0.94	0.92	0.91	0.99	0.95	1.02	1.20	1.80

The plots for the seasonal time series that was simulated is as shown below:




Fig 1: Simulated seasonal time series with noise level (σ^2) = 1

3.2 Machine Learning Techniques for forecasting

We consider 10 lag numbers as the independent variable or in other words we would like to find the association of 'y' at a particular time with its different lagged values. Lagged values are exactly the same as the original time series but they are shifted by the given time period. For example: Y_{t-1} is the same time series as Y_t but shifted back by one time period [3].

We also created lag plots to see whether the time series is random or not. We want to see whether a particular value has any correlation with its previous values. In a lag plot values Y_t are plotted against values Y_{t-k} , where 'k' represents the time period. These are the lag plots of our time series:




Fig 2: Lag plots of simulated time series

A linear lag plot indicates that the time series values are correlated to its previous time series values. We expect to see a high correlation between observations that are 12, 24 and 36 months apart due to seasonality. Therefore, we see that the lag plots of lags 12, 24 and 36 are linear. This is because the time series that was simulated was seasonal and since the frequency of the time series was 12, every 12th observation will be correlated to each other. Hence, we see that the lag 12, lag 24 and lag 36 plots are highly correlated to each other.

In the software R, we have used the embed function to gener-

ate lagged data. The dataset now has a total of 192 rows. After the formation of the dataset, we split the data into training, validation and test dataset. The last 12 values of the dataset are assigned to the test set, second last 12 values are assigned to the validation set and the rest 168 values are assigned to the training set.

In using the neural networks, the number of hidden nodes for each model vary from 2-14 with an increment of 2. Thus, the hidden layers we considered are (2,4,6,8,10,12,14). Therefore, there will be a total of 70 models for training the data.

The neuralnet function was used in order to generate the models. The hyper parameters that were tuned in the function were the number of hidden layers and the number of parameters. All the 70 models were trained on the training set. These models were then evaluated on the validation set under the criterion of root mean square error (RMSE). The best model is the one which has the lowest RMSE and this best model was then applied to the testing data. Artificial Neural networks work on a 'rolling window' approach, where we predict value that is one step ahead and then use that prediction to predict the next value. The idea is to predict $x(t+1)$, given $x(t)$ and then using $x(t+1)$, predict $x(t+2)$. In order to perform that in the software, we had to standardize the data using the scaled function before running the artificial neural network. This is to prevent larger values from dominating the result and producing the same values. The model which had the lowest RMSE when evaluated on the validation set had lags 1,2,3,4,12,13,14 and 24 as the covariates and had 8 hidden layers in it.




Fig 3: Plot of OOB error against mtry

The best mtry was automatically plugged in the random forest function in R and the ten models were trained on the training data. These 10 models were then evaluated on the validation data and the model with the lowest RMSE was selected. The best model selected also had lags 1,2,3,4,12,13,14 and 24 as the covariates. This model was then evaluated on the testing data and a final RMSE was evaluated.

4 RESULTS

The table below shows the comparison between the actual simulated time series values and the values predicted by the two machine learning algorithms: ANN and Random Forest.

Original simulated testing set	Predicted values by ANN	Predicted values by Random Forest
170.7464	169.4577	174.0609
185.3155	183.4479	183.5681
189.5513	187.7231	188.6169
211.0320	208.3834	208.9889
219.4648	217.4687	219.0205
213.3588	212.4005	216.0470
209.8160	212.9261	214.3617
233.2104	230.3436	234.2761
222.9176	222.7398	230.5797
240.9976	242.7966	253.8221
286.2513	284.3057	305.2757
425.5202	420.0849	379.2284

Looking at the values we can conclude that artificial neural network predicts values slightly closer to the original simulated time series as compared to the values predicted by the random forest algorithm.




Fig 4: Comparison between the original and Random forest values

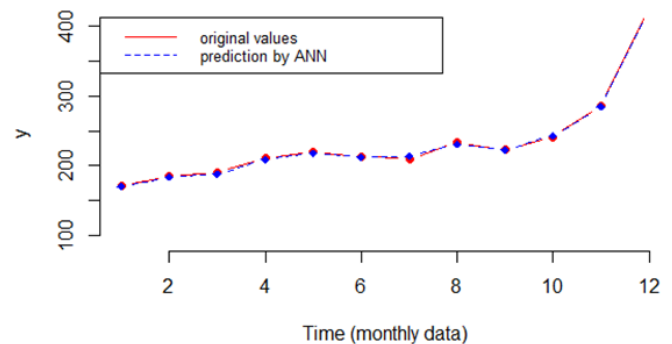


Fig 5: Comparison between the original and ANN values

Also, looking at the figures 4 and 5, we see that although the points predicted by the random forests in figure 4 lie close to the original time series values but artificial neural networks have given phenomenal results in predicting the original simulated time series as can be seen in figure 5. Thus, the predictions made by artificial neural network are more robust in comparison to the predictions made by random forest.




Fig 6: RMSE values across 100 iterations for both the algorithms

On iterating the simulations 100 times, we find that the RMSE between original Y_t values and the values predicted by the two machine learning algorithms vary fairly across a constant mean. The average RMSE given by artificial neural network is 3.07427, whereas the average RMSE between the original Y_t values and the values predicted by the random forest is 15.90248, thus giving evidence that artificial neural network is performing better than random forest.

5 CONCLUSIONS AND REMARKS

- Seasonal and trend variations are the most encountered phenomena in sectors of business and economics. In this paper, we examine the capability of two popular machine learning algorithms; artificial neural network and random forest to predict the time series with trend and seasonality.
- We found in this brief study, that for the dataset that was simulated, artificial neural network proves to be better in predicting the time series as compared to random forest.
- Future work concerns tuning other parameters of the artificial neural network like the number of decision trees, learning rate, etc.
- We shall also consider backpropagation mechanism in artificial neural network wherein the weights are adjusted to give the most accurate prediction.
- We shall also consider using a real-world time series to see which of the two methods perform better.

6 APPENDICES

The code is available in the github repository, link given in footer.

REFERENCES

- [1] Qian, X.Y. and Gao, S., 2017. Financial series prediction: Comparison between precision of time series models and machine learning methods. arXiv preprint arXiv:1706.00948.

- [2] Gupta, A. and Kumar, A., 2020, June. Mid Term Daily Load Forecasting using ARIMA, Wavelet-ARIMA and Machine Learning. In 2020 IEEE International Conference on Environment and Electrical Engineering and 2020 IEEE Industrial and Commercial Power Systems Europe (EEE-IC/I&CPS Europe) (pp. 1-5). IEEE.

- [3] Zhang, G.P. and Qi, M., 2005. Neural network forecasting for seasonal and trend time series. European journal of operational research, 160(2), pp.501-514.

- [4] Dudek, G., 2015. Short-term load forecasting using random forests. In Intelligent Systems' 2014 (pp. 821-828). Springer, Cham.

AUTHORS PROFILE

Sukhdeep Singh received Master of Data Science degree from Macquarie University, Sydney, Australia with academic merit in 2019.

Gaurangi Gupta is pursuing Bachelor's in Applied Statistics from Macquarie University and looking forward to pursue Master of Research in future.