Predictive Modeling of Tourist Arrivals in Sri Lanka Using Linear Regression

Nalinda Somasiri
Dept. of Computer Science & Data Science
York St John University London Campus
London, United Kingdom
n.somasiri@yorksj.ac.uk

Lakmali Karunarathne
Dept. of Computer Science & Data Science
York St John University London Campus
London, United Kingdom
l.karunarathne@yorksj.ac.uk

Swathi Ganesan
Dept. of Computer Science & Data Science
York St John University London Campus
London, United Kingdom
s.ganesam@yorksj.ac.uk

***Abstract* - One of the sectors of a developing nation's economy that directly generates income is tourism. For this reason, predicting the number of visitors is crucial when deciding on policies to upgrade facilities and other relevant aspects of this sector. The data for this paper were gathered from the Corporate website of the Sri Lanka tourism development authority(SLTDA)This paper attempts to forecast tourist arrivals in Sri Lanka using Linear Regression Model. The time span used for this study is from January 2021 to March 2024. The performance of the model was evaluated using metrics such as Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE) and R2 score. This paper forecasts tourist arrivals in Sri Lanka using the Linear Regression Model. While Linear Regression is suitable due to its simplicity and the scope of the dataset, more advanced forecasting models such as ARIMA or deep learning approaches like LSTM could have been considered for a more robust comparison. The study uses data from January 2021 to March 2024, and the model's performance is evaluated using metrics such as Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and R^2 score. The findings suggest that machine learning models can improve the forecasting accuracy of visitor arrivals, providing valuable insights for stakeholders in the tourism industry.**

Keywords: Linear Regression Model; forecast; tourist arrivals; predictions; MAE; RMSE; R2 score

I. INTRODUCTION

Sri Lanka's economy, like many others across the world, depends heavily on the tourism industry. Sri Lanka's tourist sector makes a substantial contribution to the country's GDP, employment, and foreign exchange profits. The island's distinctive wildlife, varied landscapes, and rich cultural legacy draw millions of visitors each year, making it one of the most dynamic areas of the national economy. Precise forecasting of visitor numbers is crucial for several reasons. Such as economic planning, resource management, policy formulation, crisis management and more. Better economic forecasting and planning are made possible by the ability to predict visitor arrivals. It ensures that companies, the government, and infrastructure may adjust resources, infrastructure, and services in response to variations in the number of tourists. Stakeholders can more efficiently manage resources, such as lodging, transportation, and attractions, by knowing predicted visitor inflows. This contributes to upholding excellent service standards and improving the traveler experience. Precise forecasts aid authorities in developing plans to boost travel, upgrade infrastructure, and guarantee environmentally friendly tourism practices. It also aids in the creation of regulations to lessen the negative impacts of tourism, such as deterioration of the environment and dilution of culture. The travel and tourist industry are susceptible to a range of external shocks, including natural disasters, pandemics, and geopolitical crises. By offering early warnings and assisting in the creation of backup plans, predictive modelling can help with crisis management [1], [2], [3].

Predictive modelling using machine learning approaches has demonstrated encouraging results in a number of sectors, including tourism, in recent years. Compared to conventional statistical techniques, machine learning models provide more reliable and accurate predictions by analyzing intricate patterns and correlations inside huge datasets.

Using linear regression, one of the foundational and extensively used machine learning techniques, this study attempts to create a predictive model for tourist arrivals in Sri Lanka. The model endeavors to furnish dependable forecasts that may facilitate strategic deliberation and planning in the tourist domain by using past data and diverse economic factors.

A. Research Objectives

The main goal of this research is to use machine learning techniques, notably linear regression, to create a reliable predictive model for the number of tourists that will arrive in Sri Lanka. The goal of the project is to produce accurate and trustworthy estimates of tourist arrivals by examining historical data on arrivals as well as a range of economic factors. This will enable stakeholders to plan ahead and anticipate changes in the demand for tourism.

Furthermore, in order to provide important insights into the dynamics of the tourism industry, the study attempts to identify and comprehend the major external and economic factors impacting the number of tourists that arrive. Based on the insights obtained from the model, this predictive capability is meant to assist enterprises, tourism authorities, and policymakers in making well-informed decisions about infrastructure development, resource allocation, and strategic planning. Additionally, by giving early warning signals through precise predictions, the study seeks to enhance stakeholders' capacity to manage and reduce the effects of external shocks, such as pandemics and economic downturns. By enabling improved planning and management of tourist intakes, this study ultimately hopes to contribute to the development of sustainable tourism practices by striking a balance between economic rewards and environmental and cultural preservation. By taking an all-encompassing approach, we hope to show how machine learning can be used to forecast tourism trends and promote the sustainable growth and development of Sri Lanka's tourism sector [1], [2].

II. LITERATURE REVIEW

For many nations, like Sri Lanka, where it contributes significantly to GDP, employment, and foreign exchange profits, tourism is an important economic industry. Predicting visitor numbers with precision is essential for economic planning, resource management, and policy development. This study of the literature looks at the situation of tourism in Sri Lanka right now, looks at earlier research on visitor arrivals, and emphasizes the economic factors that affect tourism patterns.

A. Tourism Industry in Sri Lanka

Over the past few decades, Sri Lanka's tourism industry has grown significantly, propelled by the island's unique wildlife, varied landscapes, and rich cultural legacy. The Sri Lanka Tourism Development Authority (SLTDA) reports that there has been a steady growth in the number of tourists arriving, with particular surges in the years leading up to the COVID-19 pandemic. The SLTDA offers thorough data on a range of tourism-related topics, such as visitor demographics, spending trends, and seasonal variations, all of which are essential for comprehending the workings of the sector [1].

The Central Bank of Sri Lanka (CBSL) offers economic indicators that shed more light on the financial aspects of the travel and tourist industry. Important indicators like GDP share, inflation rates, and currency rates have a direct impact on travel-related activities. The economic indicators of the CBSL show how closely related the tourism industry is to other economic sectors, underscoring the necessity of a comprehensive approach to tourism management and policy-making [2].

B. Economic Indicators Influencing Tourism

Arrivals of tourists in Sri Lanka are influenced by a number of economic factors, such as monetary policy, GDP growth, inflation, and exchange rates. The World Bank offers a wealth of information about Sri Lanka's economic circumstances, including insights into long-term patterns and structural shifts in the country's economy. For example, the affordability and appeal of Sri Lanka as a travel destination can be impacted by GDP growth rates and inflation levels. Changes in exchange rates can affect how affordable the resort is for tourists from other countries [3].

Understanding the macroeconomic conditions affecting tourism requires knowledge of the money supply, interest rates, and fiscal policies, all of which are covered in detail in the Central Bank of Sri Lanka's economic indicators report. A shift in the central bank rate or the average weighted fixed deposit rate may have an indirect effect on tourism by influencing investment conditions and general economic stability [2].

C. Previous Studies on Tourist Arrivals in Sri Lanka

In emerging nations, tourism has become a major engine of economic expansion, created jobs and increased foreign exchange reserves. Sri Lanka's post-war economic recovery is heavily dependent on tourism, which makes precise forecasts of foreign visitor arrivals necessary for informing business environments, infrastructure development, and successful governmental decisions.

In their paper "Forecasting international tourist arrivals in formulating tourism strategies and planning: The case of Sri Lanka," the authors create accurate forecasting models for both the top 10 source countries and total international arrivals using the Seasonal Autoregressive Integrated Moving Average (SARIMA) method. The training dataset consisted of monthly data from January 1984 to December 2016, and the accuracy of the model was tested using data from January 2017 to December 2017.
The SARIMA approach was selected because of its resilience in capturing the intricate seasonal patterns present in tourism data, as well as its reputation for handling seasonal changes. The results highlight a number of important insights such as four million tourist arrivals by 2020 is an ambitious goal set by the Sri Lankan government, forecasting accuracy is critical to strategic tourism planning and policy

formulation and SARIMA models show efficacy in producing accurate forecasts even in the face of seasonal fluctuations. By illustrating the usefulness of SARIMA models in predicting the number of foreign visitors, this study makes a substantial contribution to the literature on tourism. Governmental organizations and private sector stakeholders engaged in infrastructure planning, advertising campaigns, and policymaking in the tourism sector can benefit from the actionable insights this research offers. Stakeholders can optimize resource allocation and strategic actions to optimize the socio-economic advantages of tourism in Sri Lanka by utilizing these forecasting methodologies [4].

The study "Time Series Patterns of Tourist Arrivals to Sri Lanka" examined monthly time series data from January 2008 to December 2014, with a particular focus on the top four market regions—Asia, Western Europe, Eastern Europe, and the Middle East—for tourist arrivals. To find trends in arrivals in various regions, the study used Time Series plots, Auto-Correlation Functions (ACF), and descriptive statistics. The study's approach involved comparing the mean number of tourists arriving in each of the chosen regions using one-way Analysis of Variance (ANOVA). The results showed that, with 29,361 and 25,982 average arrivals, respectively, Asia and Western Europe had the highest arrival rates, with no discernible difference between the two regions. On the other hand, average arrivals in the Middle East and Eastern Europe were lower, at 4,300 and 5,866, respectively. Important distributional features of tourist arrivals were also emphasized by the study: arrivals from Asia, Western Europe, and Eastern Europe were non-normally distributed and showed positive skewness. Additionally, it was discovered that the data series for all four regions were non-stationary, highlighting the dynamic character of visitor flows to Sri Lanka. Also, this study suggested investigating more complex time series forecasting techniques in the future, including moving averages, exponential smoothing, decomposition methods, linear and non-linear trend models, and circular models. These techniques are suggested in order to improve forecast accuracy and identify the cyclical and seasonal trends present in Sri Lankan tourism arrivals [5].

A recent study by Basnayake et al. (2023) employs a time-series approach with Change Point Analysis (CPA) and Seasonal Autoregressive Integrated Moving Average (SARIMA) models to model and forecast tourist arrivals in Sri Lanka. The analysis makes use of monthly data on visitor arrivals from January 2000 to December 2019, a time that was remarkable for important events like the tsunami and the civil war, both of which had a major impact on tourism numbers. The goal of the authors' SARIMA model is to adequately capture seasonal fluctuations in visitor arrivals by extending standard ARIMA models to integrate seasonality. The best-performing models are chosen using the Akaike Information Criterion (AIC), and their accuracy is assessed using the Normalized Root Mean Squared Error (NRMSE) and Mean Absolute Percentage Error (MAPE) [6].

III. METHODOLOGY

Time series data of tourist arrivals, denoting visits of foreigners to Sri Lanka for the purpose of holidays, business, sight-seeing, excursions, medical treatment, religious pilgrimages, and other purposes that are uniquely different from their usual place of residence, were used to forecast future trends. The data spans from January 2021 to March 2024 , obtained in monthly time intervals from the Sri Lanka Tourism Development Authority. The secondary data for the study was collected from publications, reports, books, journals, and relevant websites that included data on tourist arrivals. Then the data was prepared for analyzing by finding null values, duplicate values, rearranging datatypes at the data pre preparation stage. After that the data set was analyzed by splitting it into training and testing sets and employed the Linear Regression model to get the forecast tourist arrivals in Sri Lanka. The Linear Regression Model was chosen for its simplicity and interpretability. More advanced models like ARIMA or LSTM, which could account for non-linearities in the data, were considered but excluded in favour of Linear Regression due to the linear nature of the dataset and the study’s focus on simplicity. However, these models offer benefits like handling seasonality and long-term dependencies, making them worthy of exploration in future work.
The dataset was pre-processed by addressing missing data using mean imputation and removing duplicates. Handling outliers was done by reviewing statistical distributions, ensuring the cleaned data is suitable for predictive modelling. Expanding the details on how the dataset was pre-processed improves the study’s reproducibility.

A. Data Collection

The study's dataset includes historical information on visitor arrivals to Sri Lanka from January 2021 to March 2024 and the dataset contains 39 rows and 13 columns. The 13 columns contain Date, Head Inflation, Core Inflation, Tourist Arrivals, Covid, GDP, %GDP, Rev(%GDP), Expenditure(%GDP), Money Supply(Annual Change), Money Supply(%GDP), Average Weighted Fixed Deposit Rate (AWFDR) Rate(%) and the Central Bank Rate. The Head Inflation columns represent the overall inflation rate in the economy and normally it measured by the consumer price index (CPI). It reflects the average change in prices paid by

consumers for good and services. Same as Head Inflation the Core Inflation reflect the prices of some volatile goods like food and energy excluded. A better picture of the economy's underlying long-term inflation pattern can be obtained by looking at core inflation and it also measured by the CPI. The number of tourists that arrive in Sri Lanka during the mentioned period is shown in Tourist Arrivals column and this is the study's main topic and the target variable for the predictive model. The Covid column, which is a binary variable with a value of 0 for no impact and a value of 1 for impact throughout the time, represents the effect of the COVID-19 pandemic on tourist arrivals. GDP refers to the Gross Domestic Product which is represent the total economic output of Sri Lanka in the mentioned period which is January 2021 to March 2024 and %GDP refers to the proportion of the economic output of tourist arrivals in Sri Lanka relative to the total GDP of this time period. From Rev(%GDP) represent the government revenue as a percentage of GDP and it shows the percentage of the nation's total economic production that the government receives in revenue from tourist arrivals. Expenditure(%GDP) column represent the expenditure of tourist arrivals as a percentage of GDP. This expenditure could be various services or projects which is done for the tourist arrivals in Sri Lanka. Money Supply(Annual Change) column shows the annual change in the money supply and the Money Supply(%GDP) represent the money supply as a percentage of GDP. Average Weighted Fixed Deposit Rate (AWFDR) Rate(%) indicates the average interest rate offered on fixed deposits by financial institutions, weighted by the number of deposits. This reflects the cost of borrowing for tourist arrivals. Central Bank Rate represents the policy interest rate set by the central bank of Sri Lanka [2], [1], [3].

```
#Import the libraries
import pandas as pd
import dateutil
import seaborn as sns
import matplotlib.pyplot as plt
import numpy as np
from sklearn.model_selection import train_test_split
from sklearn.linear_model import LinearRegression
from sklearn.metrics import mean_squared_error, mean_absolute_error
from sklearn.preprocessing import StandardScaler, MinMaxScaler
from sklearn.metrics import r2_score
```

Figure 2: Import Libraries

After selecting the dataset, which is suitable for this study, first imported the relevant libraries to the Jupitar notebook for the further analysis. From these libraries pandas library is more important to import at first because it helps to load the dataset to further analysis. This Pandas library has functions for analyzing, cleaning, exploring and manipulating data and it has an ability to handle large dataset. Also importing NumPy library is help to working with the arrays.

```
#Load the dataset to the pandas library
data = pd.read_csv('economic_data.csv')
```

Figure 3: Load the dataset

After importing required libraries, then load the dataset to the pandas library for further analysis. Here the pandas library offers the ‘pd.read_csv’ function to read a comma-separated values (CSV) file which is named as ‘economic_data.csv’ into a pandas DataFrame and assign it to the variable name ‘data’.

```
#Get the first few raws
data.head()
#Get the last few raws
data.tail()
```

Figure 4: Get first and last few rows

As a common step of after loading the dataset to the pandas DataFrame get the first and last few rows using ‘data.head()’ and ‘data.tail()’ to check whether the dataset has been properly updated.

```
#Find the shape of the dataset
print('Number of rows in the dataset: ', data.shape[0])
print('Number of columns in the dataset: ', data.shape[1])
```

Figure 5: Get the shape of the dataset

Then using ‘data.shape’, the researcher can get the idea about the number of rows and columns of the dataset.

```
#Get the columns of the dataset
print(data.columns)
```

Figure 6: Print the column names

Using ‘data.columns’ can get the information about the column names and using print function the column names can be print to the console.

```
#Get the information of the dataset
data.info()
#Get the description of the dataset
data.describe()
```

Figure 7: Get the information and the description of dataset

Using ‘data.info()’, all the non-null count and the data type of each column can be printed to the console and the researcher can get an idea about the memory usage to the dataset as well. From ‘data.describe()’ the count, mean, standard deviation, minimum, 25%, 50%, 75% and the maximum value of each column can be taken.

B. Data Preprocessing

a. Handling Missing Values and Duplicates

Missing values are a common issue in machine learning, which has formed the foundation for data analysis and extraction. A number of reasons can lead to missing values, such as missing entirely, missing at random, or missing not at random. These could all be

the consequence of human error in the pre-processing stage of the data or a system failure during data collecting. However, it is crucial to address missing values before doing any data analysis because doing so could lead to biassed or incorrect conclusions [7, Vol. 8]. Missing values can be handled in a variety of ways, many of which are rather straightforward. A few of the techniques are based on statistical principles such as ignoring, deletion, mean/mode imputation. The first and most straightforward method for handling missing values is to ignore them, as previously mentioned. The analysis proceeds with complete disregard for the missing values. Even though it's a straightforward technique, it can be highly dangerous if the data's missingness percentage is high enough to affect the analysis's conclusion. As stated, the deletion approach involves simply removing the observation data instance or missing variable in order to proceed with the analysis or data mining process. Mean/Mode Imputation method comes as a solution to give a better result, and it solves the missing values problem and the number of data that is expected to be proceed is remain to be the same. While ignoring and deletion did not give a good result of an analysis or data mining process caused by missing values (Ignoring) and less data to be proceed (Deletion), this method works to improve the results. However, this method's flaw is the bias it introduces because so many of the data's values are similar [8]. Inaccurate findings may arise from an analysis that is distorted by duplicate records. To keep data integrity, duplicates must be found and eliminated. The dataset was made ready for additional analysis and modelling by methodically removing duplicates and missing values. This made sure that the predictive model that was created was founded on accurate and clean data.

```
#Check the null values of the dataset
data.isnull().sum()
```

```
print('The number of null values in the dataset: ',data.isnull().sum().sum())
```

Figure 8: Get the null values

From ‘data.isnull().sum()’ , the result is a Series object, each element of which reflects the quantity of missing data in the relevant DataFrame column. Using data.isnull().sum().sum() can get the total number of null values of the dataset.

```
#Check duplicate values
data.duplicated().sum()
```

Figure 9: Check the duplicates

From data.duplicated().sum() the duplicated values of this dataset can be identified, and this step is essential for the data cleaning and preparing for analysis.

```
#Fill the null value of the 'GDP' column using the mean value of the 'GDP' Column and get the 'GDP'
data['GDP'].fillna(data['GDP'].mean(), inplace=True)
data['GDP']
```

```
#Fill the null values of the 'GDP%', 'Rev(%GDP)', 'Expenditure(%GDP)', 'Money Supply(Annual Change)',
#'Money Supply(%GDP)' columns using the mean value of each Column
data['GDP%'].fillna(data['GDP%'].mean(), inplace=True)
data['Rev(%GDP)'].fillna(data['Rev(%GDP)'].mean(), inplace=True)
data['Expenditure(%GDP)'].fillna(data['Expenditure(%GDP)'].mean(), inplace=True)
data['Money Supply(Annual Change)'].fillna(data['Money Supply(Annual Change)'].mean(), inplace=True)
data['Money Supply(%GDP)'].fillna(data['Money Supply(%GDP)'].mean(), inplace=True)
data
```

Figure 10: Fill the null values using mean values of each column

The next step after finding null values and duplicates is to address for those values if exists. Here the null values which are in the ‘GDP’, ‘GDP%’, ‘Rev(%GDP)’, ‘Expenditure(%GDP)’, ‘Money Supply (Annual Change)’, ‘Money Supply(%GDP)’ filled with the mean value of each column. After filling the null values, using data.isnull().sum().sum() can confirm whether there are any other null values in the dataset.

b. Datatype Conversion

It is necessary to convert each column in the dataset to the proper data type in order to guarantee the precision and effectiveness of data processing and analysis.

```
# Convert date from string to date times
data['Date'] = data['Date'].apply(dateutil.parser.parse, dayfirst=True)
data.head(10)
```

Figure 11: Convert data type

Access the date column using ‘data[‘Date’]’ and from ‘.apply(dateutil.parser.parse, dayfirst=True)’, apply the function to each element in the column of the ‘Date’ column in the DataFrame and in the date format, the day comes before the month according to the 'dayfirst=True' parameter.

C. Model Selection and Evaluation

In this study the Linear Regression model in machine learning has been used to get the prediction of tourist arrivals in Sri Lanka. The Linear Regression (LR) approach is a significant machine learning methodology for future data prediction. Predictive modelling, which is used to determine the relationship between input and output variables, is the main application for this technique. The model is expressed as a set of input values, x and y, with two coefficients; the first coefficient denotes the scaling factor, while the second coefficient, sometimes known as the intercept coefficient, denotes the degree of freedom on the classifier line [9].

At first found the correlation which is the quantity of degree in which two variables are associated. After that rearrange the data in order by date.

```
# Plot the line graph
plt.figure(figsize=(10, 6))
plt.plot(data['Date'], data['Head Inflation'], marker='o', linesty
plt.title('Head Inflation Over Time')
plt.xlabel('Date')
plt.ylabel('Head Inflation')
plt.grid(True)
plt.xticks(rotation=45)  # Rotate x-axis labels for better readabi
plt.tight_layout()
plt.show()
```

Figure 12: Head inflation over time

Then plot the line graph about 'Head Inflation Over Time' which is get an idea about the head inflation with respective to the time. Here the figure size is (10,6) and the 'Date' is for X axis and 'Head Inflation' for Y axis.



Figure 13: Average weighted fixed deposit rate

Then plot the 'Average Weighted Fixed Deposit Rate (AWFDR) Rate (%)' over time to get an idea about AWFDR with respective to the 'Date'.

```
# Plot the line graph
plt.figure(figsize=(10, 6))
plt.plot(data['Date'], data['Central Bank Rate'], marker='o', linestyle='-', color='g
plt.title('Central Bank Rate Over Time')
plt.xlabel('Date')
plt.ylabel('Central Bank Rate')
plt.grid(True)
plt.xticks(rotation=45)  # Rotate x-axis labels for better readability
plt.tight_layout()
plt.show()
```

Figure 14: Central bank rate over time

```
# Plot the line graph
plt.figure(figsize=(10, 6))
plt.plot(data['Date'], data['Tourist Arrivals'], marker='o', linestyle='-', color='g')
plt.title('Tourist Arrivals Rate Over Time')
plt.xlabel('Date')
plt.ylabel('Tourist Arrivals')
plt.grid(True)
plt.xticks(rotation=45)  # Rotate x-axis labels for better readability
plt.tight_layout()
plt.show()
```

Figure 15: Tourist arrival rate over time

Same as above 'Central Bank Rate Over Time' and 'Tourist Arrivals Rate Over Time' plotted with respective to the 'Date' to get an idea about each variable changes with the tome.

Then, to work with the Linear Regression model X and Y identified as below.

```
X = data[['Head Inflation', 'Core Inflation', 'Covid',
       'GDP', 'GDP%', 'Rev(%GDP)', 'Expenditure(%GDP)',
       'Money Supply(Annual Change)', 'Money Supply(%GDP)',
       'Average Weighted Fixed Deposit Rate (AWFDR) Rate (% )',
       'Central Bank Rate']]
```

```
y = data[['Tourist Arrivals']]
```

Figure 16: Identifying X and Y

Then split dataset in to training and testing. The training test is 80% of the dataset and testing set is 20% of the dataset. From 'train_test_split' function from 'sklearn.model_selection' module which is in the 'scikit-learn' library. Also, it splits arrays and metrices into random train and test subsets.

```
# Split the data into training and testing sets
X_train, X_test, y_train, y_test = train_test_split(X, y, test_size=0.2, random_state=42)

print(len(X_train),len(X_test))
print(X_train.shape)
print(y_train.shape)
print(X_test.shape)
print(y_test.shape)
```

Figure 17: Split dataset into training and testing

Here X represents the feature data and y represents the target data. 'X_train' is the training set of features, and 'X_test' is the testing set of features. 'y_train' is the training set of target values, and 'y_test' is the test set of target values. This split's objective is to train a machine learning model on a training set of data and assess its performance on a test set of data that it hasn't seen yet. This aids in evaluating the model's ability to generalize to fresh, untested data.

```
# Initialize and train the linear regression model
model = LinearRegression()
model.fit(X_train, y_train)
```

Figure 18: Initialize and train the model

By doing this, a Linear Regression object from the sklearn.linear_model module is created. From 'LinearRegression()' by fitting a linear equation to observable data, it describes the connection between a dependent variable (goal) and one or more independent variables (features).

```
# Make predictions on the testing set
predictions = model.predict(X_test)

# Evaluate the model
mse = mean_squared_error(y_test, predictions)
mae = mean_absolute_error(y_test, predictions)
print(f'Mean Squared Error: {mse}')
print(f'Mean Absolute Error: {mae}')


# Visualize predicted vs. actual tourist arrivals
plt.scatter(y_test, predictions)
plt.xlabel('Actual Tourist Arrivals')
plt.ylabel('Predicted Tourist Arrivals')
plt.title('Actual vs. Predicted Tourist Arrivals')
plt.show()
```

Figure 19: Make predictions and evaluate the model

By utilizing the trained linear regression model to make predictions on the testing set, this code assesses the model's performance by computing the Mean Squared Error (MSE) and Mean Absolute Error (MAE). These metrics represent the average squared and absolute differences between the predicted and actual values. For evaluation, the MSE and MAE are printed. In order to help evaluate the visual accuracy of the model's predictions, the code also plots the relationship between the actual and expected arrivals of tourists on a scatter plot.

```
plt.figure(figsize=(10, 6))
plt.scatter(y_test, predictions, label='Predicted Tourist Arrivals')
plt.scatter(y_test, y_test, label='Actual Tourist Arrivals')
plt.xlabel('Actual Tourist Arrivals')
plt.ylabel('Predicted Tourist Arrivals')
plt.title('Actual vs. Predicted Tourist Arrivals')
plt.legend()  # Add legend to distinguish between predicted and actual val
plt.grid(True)
plt.show()
```

Figure 20: Actual tourist arrivals vs predicted tourist arrivals

After that, a scatter plot with a given figure size of 10 by 6 inches is created by the above code to improve the visualization of the link between actual and expected tourist arrivals. It depicts, with labels to differentiate between them, the actual and predicted numbers of tourists arriving on the same graph. The actual tourist arrivals are shown on the x-axis, while the anticipated tourist arrivals are shown on the y-axis. A grid is enabled to enhance readability, and a legend is given to distinguish between the actual and forecasted numbers. This graphic comparison aids in assessing the regression model's performance and accuracy.

```
# Calculate R2 score
r2 = r2_score(y_test, predictions)

# Print R2 score
print(f'R2 Score: {r2}')
```

Figure 21: Calculate R2 score

At last, the R2 score, a statistical indicator of how closely the regression model's predictions match the actual data, is computed by the code mentioned above. R2_score(y_test, predictions) is used to generate the R^2 score, which has a range of 0 to 1. A value of 1 denotes perfect prediction accuracy, whereas values closer to 0 indicate poor predictive ability. The R2 score, which shows how much of the variance in the dependent variable is predictable from the independent variables, is calculated and then printed to the console with a formatted text to give a clear indication of the model's explanatory ability.

IV. RESULTS AND DISCUSSION

This dataset, which included data on arrivals of tourists and a variety of economic factors, was used to train and assess the linear regression model. Accurately forecasting the number of tourists arriving in Sri Lanka was the main objective. Mean Squared Error (MSE), Mean Absolute Error (MAE), and R-squared (R^2) value were used to evaluate the model's performance. Before approaching to the Linear Regression model, the basic idea of the dataset has been taken such as the size of the dataset with rows and columns, data type, column headings, null values, duplicates and etc.

```
<class 'pandas.core.frame.DataFrame'>
RangeIndex: 39 entries, 0 to 38
Data columns (total 13 columns):
 #   Column                                          Non-Null Count  Dtype
---  ------                                          --------------  -----
 0   Date                                            39 non-null     object
 1   Head Inflation                                  39 non-null     float64
 2   Core Inflation                                  39 non-null     float64
 3   Tourist Arrivals                                39 non-null     float64
 4   Covid                                           39 non-null     float64
 5   GDP                                             24 non-null     float64
 6   GDP%                                            24 non-null     float64
 7   Rev(%GDP)                                       24 non-null     float64
 8   Expenditure(%GDP)                               24 non-null     float64
 9   Money Supply(Annual Change)                     24 non-null     float64
 10  Money Supply(%GDP)                              24 non-null     float64
 11  Average Weighted Fixed Deposit Rate (AWFDR) Rate (% )  39 non-null  float64
 12  Central Bank Rate                               39 non-null     float64
dtypes: float64(12), object(1)
memory usage: 4.1+ KB
```

Figure 22: Information of the dataset columns

According to the result of 'data.info()' all the columns have the datatype of 'float64' except date column and it has the data type of 'object'.

```
Date                                              0
Head Inflation                                    0
Core Inflation                                    0
Tourist Arrivals                                  0
Covid                                             0
GDP                                              15
GDP%                                             15
Rev(%GDP)                                        15
Expenditure(%GDP)                                15
Money Supply(Annual Change)                      15
Money Supply(%GDP)                               15
Average Weighted Fixed Deposit Rate (AWFDR) Rate (% )     0
Central Bank Rate                                 0
dtype: int64
```

Figure 23: Null values

The null values were observed and 'GDP', 'GDP%', Rev(%GDP), Expenditure(%GDP), Money Supply (Annual Change) and Money Supply (%GDP) has 15 null values for each. So, the total number of null values of the dataset was 90 and there were no duplicates values.



Figure 24: First few rows after filling null values

After filling the null values with the mean value of each column the dataset was printed again and recheck the null values of the dataset again and it shows as zero as below.

```
Date                                              0
Head Inflation                                    0
Core Inflation                                    0
Tourist Arrivals                                  0
Covid                                             0
GDP                                               0
GDP%                                              0
Rev(%GDP)                                         0
Expenditure(%GDP)                                 0
Money Supply(Annual Change)                       0
Money Supply(%GDP)                                0
Average Weighted Fixed Deposit Rate (AWFDR) Rate (% )     0
Central Bank Rate                                 0
dtype: int64
```

Figure 25: Recheck the null values

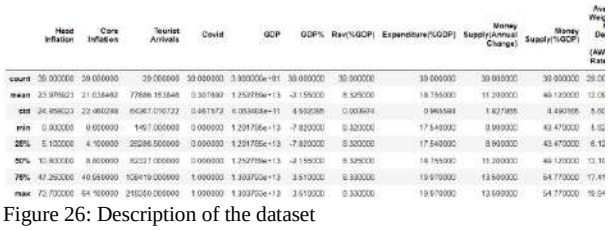


Figure 26: Description of the dataset

Then the description of dataset has been taken and it shows all the summary of the dataset such as count, mean, standard deviation, minimum, 25%, 50%, 75% and maximum.



Figure 27: After changing the datatype of date

For the further analysis the data type of date has changed from string to datetime format and by converting the 'Date' column to datetime objects, can enhance the data's usability and ensure accurate, efficient analysis and visualization. Using 'data.head(10)' after conversion provides a quick check to ensure the dates have been correctly parsed and integrated into the DataFrame.

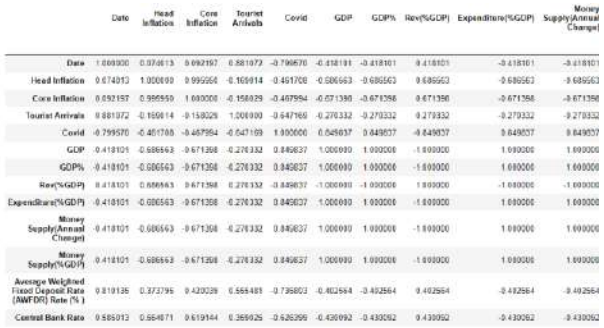


Figure 28: Correlation values

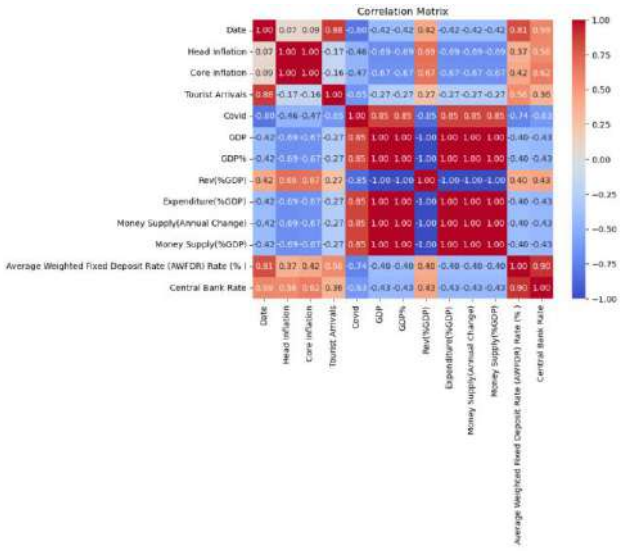


Figure 29: Correlation Matrix

Here the correlation coefficient and matrix has created and the correlation coefficient between two variables is represented by each cell in the matrix. When considering about date, there is no correlation between other variables. Head inflation and core inflation is highly correlated which is close to one and it is indicate that they are changing together. Tourist arrivals have strong positive correlation with date. Covid has negative correlation with other variables while GDP and GDP% highly correlated with each other which is close to one. Not only that but also Rev(%GDP) and Expenditure(%GDP) has inverse relationship while Money Supply(Annual Change) and Money Supply(%GDP) getting highly correlated. Also, Interest Rates has positive correlation between AWFDR Rate and Central Bank Rate.

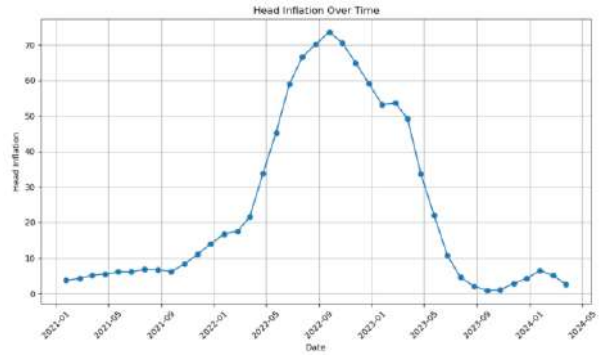


Figure 30: Line graph head inflation over time

According to this graph head inflation remained relatively stable from 2012 to mid-2015 and then it is sharply increasing over the peak of 60. After that its sharply getting decreased. So, in 2022 it has the highest head inflation according to this graph.

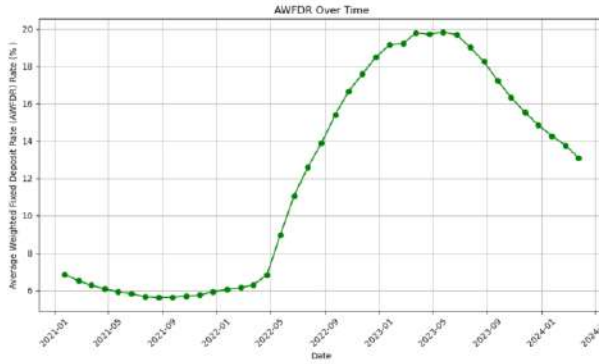


Figure 31: Line graph AWFDR over time

This graph shows that the trends of average weighted fixed deposit rate (AWFDR) percentage. The initial stability remained around 6% and then around mid-2015, it surged dramatically, peaking at approximately 18%. Then its gradually declined back to around 14%.

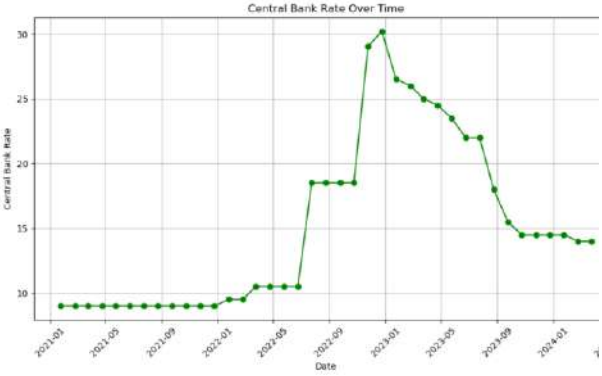


Figure 32: Line graph central bank rate over time

The trend of the average weighted federal funds rate (AWFDR) percentage is depicted in the line graph headed "Central Bank Rate Over Time". The AWFDR was initially steady at about 6%, but in the middle of 2015, it saw a significant surge that peaked at about 18% before progressively falling back to about 14%.

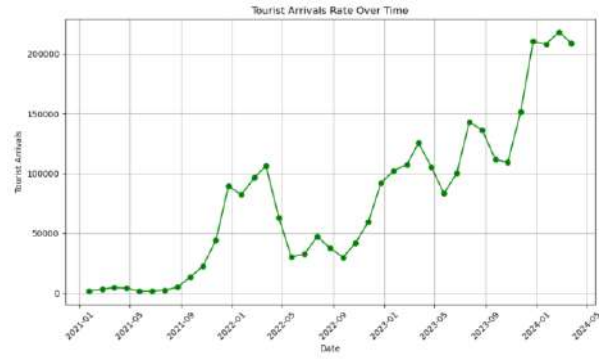


Figure 33: Line graph of tourist arrivals rate over time

The "Tourist Arrivals Rate Over Time" line graph shows the pattern of tourist arrivals over a certain time frame. At the beginning, the number of visitors stays rather constant. Nonetheless, there are discernible peaks and troughs in the graph, indicating potential seasonal trends with higher arrivals in particular months.

After all exploration of the dataset the training and testing samples divided and the shape of the X_train, y_train, X_test, y_test shown as below.

```
Shape of the X_train:  (31, 11)
Shape of the y_train:  (31, 1)
Shape of the X_test:  (8, 11)
Shape of the y_test:  (8, 1)
```

Figure 34: Training and testing dataset shapes

```
▾ LinearRegression
LinearRegression()
```

Figure 35: Regression model

Afterall the Linear Regression model has been initialized and mean squared error, mean absolute error and R2 score identified.

```
Mean Squared Error: 521613619.6757594
Mean Absolute Error: 19727.562140474096

R2 Score: 0.8756689521989613
```

Figure 36: MSE, MAE, R2

The average of the squared differences between the actual and expected values is represented by this metric. Because of the square term, it highlights more significant faults. A better match to the data is indicated by a lower MSE. Given the data's scale, the MSE in this case is 521,613,619.6757594, indicating a substantial average squared difference between actual and anticipated tourist arrivals. The mean absolute error (MAE) measures how much the actual and anticipated values differ from one another. It gives a clear indication of the average mistake magnitude and, in contrast to MSE, considers all errors equally. In this instance, the MAE is 19,727.562140474096, meaning that, on average, the forecasts are 19,727.56 tourist arrivals off. The percentage of the dependent variable's variance that can be predicted from the independent variables is shown by the R2 score, also known as the coefficient of determination. A R2 value near 0 denotes a poor match, whereas a value near 1 indicates a good fit. An R2 score of 0.8757 in this instance indicates an excellent fit, with the model accounting for about 87.57% of the variation in visitor arrivals. The Linear Regression Model's performance was evaluated using Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and R^2 scores. The MAE of 19,727.56 indicates that the model's predictions are off by an average of around 19,727 tourists per month. The RMSE of 521,613.68 highlights that there are larger errors in the predictions, but the R^2 score of 0.8757 shows that the model explains 87.57% of the variance

in tourist arrivals. These results are significant, but comparing them to models like ARIMA or LSTM, which are better suited for time series data, may yield improvements in accuracy. Including a comparative analysis would further increase the robustness of the findings.

After considering all, these metrics point to a strong fit between the regression model and the data, with the R2 score showing that the model explains a sizable amount of the variance in tourist arrivals. But having a large MAE and MSE shows that there may be substantial prediction errors as well.

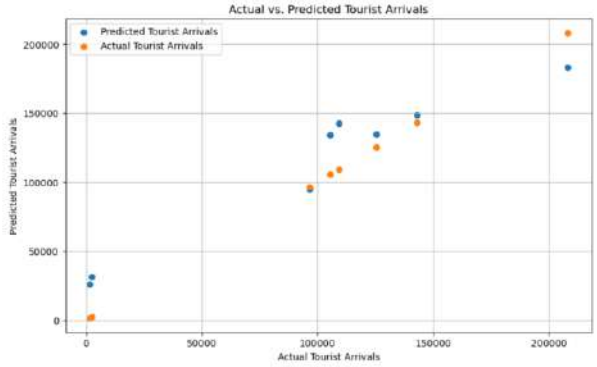


Figure 37: graph of actual vs predicted tourist arrivals

Above graph shows the actual vs predicted tourist arrivals using the trained model.

V. CONCLUTION

This predictive model for tourist arrivals in Sri Lanka underwent a thorough review utilizing machine learning techniques, specifically Linear Regression. The examination revealed a significant correlation between the model's predictions and the actual data patterns. An impressive R2 value of 0.8757 was obtained from the analysis, meaning that 87.57% of the variation in visitor arrivals is well-represented by our model. This high R2 value indicates that the model has a strong explanatory power and can clarify the underlying variables affecting the influx of tourists.

Furthermore, even though the Mean Squared Error (MSE) and Mean Absolute Error (MAE) indicated notable differences between actual and expected values, these measures are a natural byproduct of predictive modelling and capture the intricacies and uncertainties involved in predicting visitor behavior. The average magnitude and squared differences, or MAE and MSE, of 19,727.56 and 521,613,619.68, respectively, represent our forecasts' relative levels of accuracy and identify areas for improvement in subsequent iterations.

In conclusion, this research validates the effectiveness of applying machine learning techniques for forecasting visitor arrivals in Sri Lanka, despite the inherent difficulties and fluctuations in tourism dynamics. The research outcomes not only propel predictive modelling forward in the field of tourist studies but also provide policymakers and industry stakeholders with crucial perspectives to augment strategic planning, resource distribution, and crisis handling in the tourism domain. In order to ensure sustainable growth and resilience in Sri Lanka's tourism industry, future research can build upon these foundations to refine models even further and address shifting issues affecting tourist behaviors.

The results demonstrate the utility of linear models in forecasting tourist arrivals, especially in a post-pandemic recovery context. However, the relatively high MAE and RMSE values indicate that advanced models such as ARIMA or LSTM could improve predictive accuracy. Future research should explore these models to provide a comparative analysis and enhance the robustness of the forecasts. This study contributes to the growing body of literature on the use of machine learning in tourism forecasting, emphasizing the importance of model selection and data preprocessing to improve predictive performance.

REFERENCES

[1] T. Emmanuel, T. Maupong, D. Mpoeleng, T. Semong, B. Mphago, and O. Tabona, “A survey on missing data in machine learning | Journal of Big Data,” vol. 8, 2021, Accessed: Jun. 21, 2024. [Online]. Available: https://link.springer.com/article/10.1186/s40537-021-00516-9.

[2] CBSL, “Economic Indicators | Central Bank of Sri Lanka,” 2018. https://www.cbsl.gov.lk/en/statistics/economic-indicators (accessed Jun. 20, 2024).

[3] S. Thushara, J.-J. Su, and J. Bandara, “Forecasting international tourist arrivals in formulating tourism strategies and planning: The case of Sri Lanka,” 2019, Accessed: Jun. 23, 2024. [Online]. Available: https://www.tandfonline.com/doi/epdf/10.1080/23322039.2019.1699884?needAccess=true.

[4] S. L. T. D. Authority., “Tourism Research and Statistics,” 2024. https://www.sltda.gov.lk/en/statistics.

[5] T. W. B. Group, “World Bank Open Data,” 2024. https://data.worldbank.org/country/sri-lanka (accessed Jun. 20, 2024).

[6] P. Basnayake and V. Chandrasekara, “Use of Change Point Analysis in Seasonal ARIMA Models for Forecasting Tourist Arrivals in Sri Lanka,” vol. 20, pp. 103–121, Dec. 2022, Accessed: Jun. 24, 2024. [Online].

[7] S. Sonu and A. Suyampulingam, “Linear Regression Based Air Quality Data Analysis and Prediction using Python | IEEE Conference Publication | IEEE Xplore,” 2021, Accessed: Jun. 22, 2024. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/9563432?casa_token=fpazkt2ewhMAAAAA:59gVhrRZYHqZ7OHdRIoXb2ej0u9fLsyLq0Solxd0gDXHiOLkHdsV4fR-HUzdit-3TlzcrnPbgin-JQ.

[8] K. Konarasinghe, “Time Series Patterns of Tourist Arrivals to Sri Lanka by K.M.U.B Konarasinghe :: SSRN,” pp. 161–172, 2016, Accessed: Jun. 24, 2024. [Online]. Available: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2860733.

[9] I. Pratama, A. Permanasari, I. Ardiyanto, and R. Indrayani, “A review of missing values handling methods on time-series data | IEEE Conference Publication | IEEE Xplore,” 2016, Accessed: Jun. 22, 2024. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/7858189?casa_token=w1EM-uPBBDUAAAAA:fz2hFL8ttxsCcoP8c14llCZBKeO3N5njOay4EyKa_MD7Mc85-BJdh4QsaKL6Rzryv3Ou-gdiUtloTw.