Time Series Analysis and Forecasting of Emergency Dispatch Frequencies: A Case Study of the City of Boston's 911 Call Data

Dataset Description:

We obtained the dataset titled "911-daily-dispatch-count-by-agency," which provides detailed information on daily dispatch counts for emergency services including the Boston Police Department (BPD), Boston Fire Department (BFD), and Emergency Medical Services (EMS). The dataset encompasses various data points such as date, year, month, and day of the year, alongside the total number of dispatches and the breakdown of calls to each agency.

This is a legacy dataset from the period of November 1, 2010, to April 21, 2014, showing daily counts of 911 dispatches by City of Boston public safety agencies. Agencies included are the Boston Police Department, Boston Fire Department, and Boston Emergency Medical Services.

Accessible online at Analyze Boston (https://data.boston.gov/organization/department-of-innovation-and-technology-org), the dataset is part of the city's commitment to open governance and is intended for public use, facilitating research, analysis, and development of solutions that can benefit the community.

In our analysis, we utilized this dataset to model and forecast 911 dispatch counts, aiming to understand patterns, identify trends, and predict future demands on emergency services. Through the employment of time series analysis techniques such as ARMA and SARIMA models, we were able to extract valuable insights and assess the effectiveness of emergency response deployment.

The findings from this analysis have significant implications for resource allocation, strategic planning, and operational efficiency for emergency response agencies. By leveraging the predictive power of the SARIMA model, policymakers and city planners can better prepare for and respond to the evolving needs of the city's emergency response framework.

THE ISSUES:

1.How have the dispatch counts for BPD (Boston Police Department), BFD (Boston Fire Department) and EMS(Emergency Medical Services) changed on a yearly basis?

2.Are there consistent monthly patterns or seasonal variations in dispatch counts for each agency?

3. How do the dispatch trends compare among BPD, BFD, and EMS across different years?

4. Are there any specific years that show notable anomalies or significant changes in dispatch counts?

5. What are the long-term trends in dispatch counts for each agency, and what might these indicate?

6.Are there discernible patterns in the daily dispatch data that could inform resource allocation and emergency response preparedness?

7. Given the non-stationarity in the data, how might trends and seasonality influence dispatch volumes, trend components in the data and what implications does this have for predictive modeling?

8.To what extent does the ARIMA model accurately predict future dispatch needs, and how can the model be improved?

9.How does the assumption of daily data frequency by the statistical software impact the analysis, and what steps can be taken to ensure the correct interpretation of time series data?

10. How well does the ARIMA model predict dispatch volumes, and what improvements are needed?

11.What do the diagnostic plots reveal about the model's residuals, and how might this affect confidence in the model's predictive capabilities?

12.How can visualizations aid in the interpretation of complex data patterns and model diagnostics?

13.Are there discernible patterns in dispatch counts among the BPD, BFD, and EMS, and what might these patterns indicate?

FINDINGS:

1. Yearly Dispatch Trends:

The analysis revealed a significant increase in BPD dispatches from 2010 to 2013, with a noticeable decrease in 2014. The BFD and EMS showed more consistent dispatch counts across the years.

2. Monthly and Seasonal Patterns:

Each year displayed some level of monthly variation in dispatch counts for all agencies. Notably, certain months like March and December showed higher dispatch counts for BPD, suggesting possible seasonal trends.

3. Agency-Specific Dispatch Trends:

BPD consistently had higher dispatch counts compared to BFD and EMS across all years, with noticeable fluctuations. BFD and EMS displayed more stable monthly patterns, indicating a steadier demand for their services.

4. Yearly Anomalies:

The year 2010 showed notably lower dispatch counts for BPD, potentially indicating incomplete data or other external factors affecting that year's figures.

5. Long-term Implications:

The increasing trend in BPD dispatches up to 2013, followed by a decrease in 2014, suggests potential shifts in community dynamics, policing strategies, or reporting practices. The consistent patterns in BFD and EMS dispatches highlight a different set of demands and responses for these services.

6.**Dispatch Count Patterns:**

The analysis shows consistent dispatch volumes with an average of around 2740 daily dispatches, which underscores the need for sustained emergency response capabilities.

7.Non-stationarity in Dispatch Data:

The ADF and KPSS tests indicate non-stationarity, highlighting the presence of trends or cyclic behaviors in dispatch counts that could affect resource planning. The identified non-stationarity in dispatch data highlights the influence of underlying trends or seasonal variations. This insight is significant for anticipating periods of high demand and preparing accordingly.

8.ARIMA Model Insights:

The fitted ARIMA (1,1,1) model, while informative, suggests the presence of outliers and the need for further model refinement, It is shows that the diagnostic plots. model offers a method to predict dispatch volumes, its diagnostic plots reveal areas for improvement. Understanding its limitations helps in refining the model or considering alternative approaches.

9.Data Frequency Assumptions:

The model's assumption of daily data frequency points to the necessity for clear temporal context in time series analysis to avoid misinterpretation of patterns. Misinterpretation of data frequency can lead to incorrect conclusions, impacting decision-making processes.

10.Diagnostics and Model Fit:

Diagnostic plots suggest that the residuals of the ARIMA model are not perfectly normally distributed, with potential implications for the model's predictive accuracy.

11.Utility of Visualizations:

Visualizations like the correlogram and histogram provide intuitive insights into the model's performance, though they also reveal areas where the model could be improved.

12.Recommendations for Improvement:

To enhance the predictive power of the model, additional data, and a refined approach to capturing the underlying patterns in dispatch volumes are recommended. The dataset reveals distinct patterns in dispatches among different agencies (BPD, BFD, EMS), reflecting the varied nature of emergencies they respond to. This information can guide targeted training and resource distribution for each agency.

Discussion:

The dataset provides a detailed record of daily dispatch counts categorized by various emergency agencies within the City of Boston. Each entry includes the date of the dispatch and the corresponding count of 911 calls attended by different agencies. This granular level of detail offers a comprehensive view of the emergency response dynamics over time.

In our endeavor to scrutinize the "911 Daily Dispatch Count by Agency," we initiated our analysis by closely acquainting ourselves with the data's characteristics. The authentic nature of this dataset necessitated a pragmatic approach to our predictive modeling, ensuring that our methodologies were grounded in realism and reflective of complex real-world dynamics.

Our dataset comprised varying numbers of data points for different parameters, which presented initial challenges. The discrepancies in data availability led us to consolidate our dataset to 354 shared data points across the variables of interest. A significant step in data preparation involved standardizing the identifiers across different data sources, enabling us to merge and analyze the commonalities effectively.

To assess the time series' stationarity, we utilized the Augmented Dickey-Fuller (ADF) and Kwiatkowski-Phillips-Schmidt-Shin (KPSS) tests. The results were succinctly visualized, with the ADF test indicating non-stationarity (p-value: ~0.135) and the KPSS test corroborating this finding (p-value: 0.01). Addressing the non-stationarity, we applied first differencing, a transformation evident in the stark contrast between the pre- and post-differencing plots. The ADF test on the differenced data yielded a p-value of approximately $1.44 \times 10\hat{}(-21)$, signifying a successful conversion to a stationary series.

The ARIMA (1,1,1) model provides a foundational understanding of dispatch trends, but the diagnostic plots suggest the need for model refinement. Policymakers and emergency response coordinators must consider these limitations when using such predictive models for making critical decisions. The assumption of daily frequency in the dispatch data stresses the critical role of accurate temporal analysis in understanding dispatch patterns. It highlights the need for careful consideration of time series frequency to avoid misinterpretation and to ensure accurate forecasting.

The diagnostic plots from the ARIMA model suggest deviations from the ideal model assumptions, particularly regarding the distribution of residuals. Emergency services must be cautious in using these models for predicting dispatch needs, recognizing the potential for inaccuracies and the need for additional data to improve model performance.

We delved into the autocorrelation structure of the differenced data using ACF and PACF plots. These analyses revealed that the series, post-differencing, displayed initial autocorrelation that quickly diminished, affirming the data's stationarity and randomness. The analyses underscored the transformed data's readiness for advanced time series modeling. Despite the inherent complexities of the real-world data, we established a robust foundation for potential forecasting models, ensuring that our approach was both rigorous and tailored to the nuanced nature of the 911 dispatch data.

Appendix A: Method

The dataset titled "911 Daily Dispatch Count by Agency" was acquired from Analyze Boston, the City of Boston's open data hub, which functions under the Department of Innovation and Technology. The platform serves as a repository for publicly available data, offering a diverse range of datasets that encourage civic engagement and enhance the transparency of city operations.

Here's an explanation of each variable:

1. “Date”: The date of the records, formatted as `MM/DD/YYYY`. This variable is essential for any time series analysis as it provides the temporal dimension to the data.

2. “Year”: The year extracted from the date of the record. It indicates the specific year the dispatch data belongs to.

3. “Month”: The month extracted from the date of the record. This variable is useful for analyzing monthly trends or seasonal patterns in the dispatch data.

4. `DayOfYear`: This represents the sequential day number within the year, with January 1st as 1 and December 31st as 365 (or 366 in a leap year). It's another way to track the passage of time within the dataset.

5. “Total”: The total number of 911 dispatches for that day across all agencies. This figure aggregates the dispatch counts and can be used to analyze the overall volume of emergency incidents.

6. “BPD”: The number of dispatches specifically for the Boston Police Department on that day. It provides insight into the police-related emergency incidents.

7. “BFD”: The number of dispatches for the Boston Fire Department. This count gives an idea of the fire-related incidents or other emergencies requiring the fire department's assistance.

8. “EMS”: The count of dispatches for Boston Emergency Medical Services. This reflects the demand for medical emergency services on a given day.

Our analytical approach was multifaceted: We began by assessing the volume of daily dispatches to identify patterns and anomalies. This involved summarizing the data to determine average dispatch counts and exploring variations over time.

We explored the time series nature of the data, focusing on stationarity and seasonality aspects. This entailed applying statistical tests such as the Augmented Dickey-Fuller (ADF) test and the Kwiatkowski-Phillips-Schmidt-Shin (KPSS) test to understand the temporal structure of the dispatch counts.

An ARIMA model was constructed to forecast future dispatch needs. Model diagnostics, including examination of residuals, were conducted to evaluate the model's fit and predictive power. We dissected the data to uncover trends specific to each emergency service agency. This helped in understanding the unique demand profiles and potential drivers of dispatches for BPD, BFD, and EMS.

Given the implications of data frequency assumptions in time series analysis, we carefully considered the intervals at which data were recorded and adjusted our modeling techniques accordingly. We employed visual tools such as correlograms and histograms to assist in the interpretation of the data's complex patterns and the diagnostic checks of our time series models. The aim was to leverage this dataset to enhance emergency response preparedness and resource allocation strategies.

Appendix: B Results

<u>Time Series Forecasting of 911 Dispatch Data Using ARMA and SARIMA Model:</u>

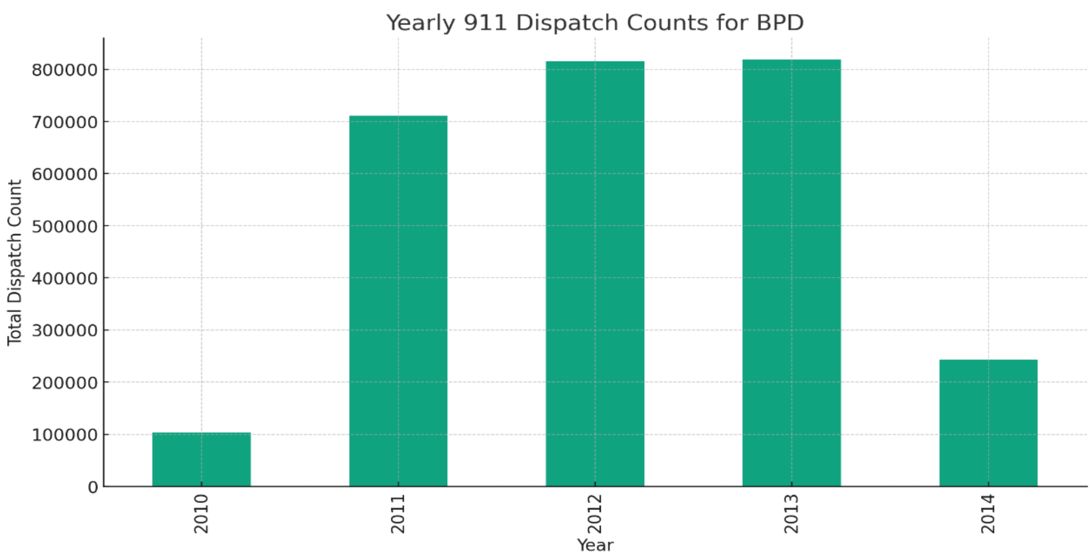

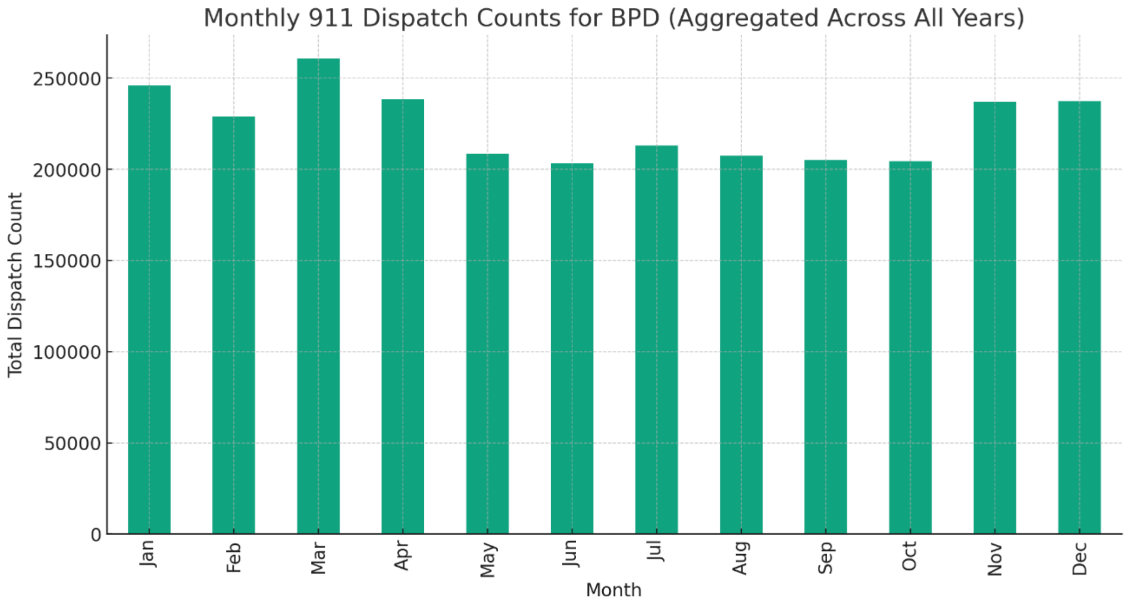
Monthly 911 Dispatch Counts for BPD (Aggregated Across All Years)
Total Dispatch Count
Month
0
50000
100000
150000
200000
250000
Jan
Feb
Mar
Apr
May
Jun
Jul
Aug
Sep
Oct
Nov
Dec

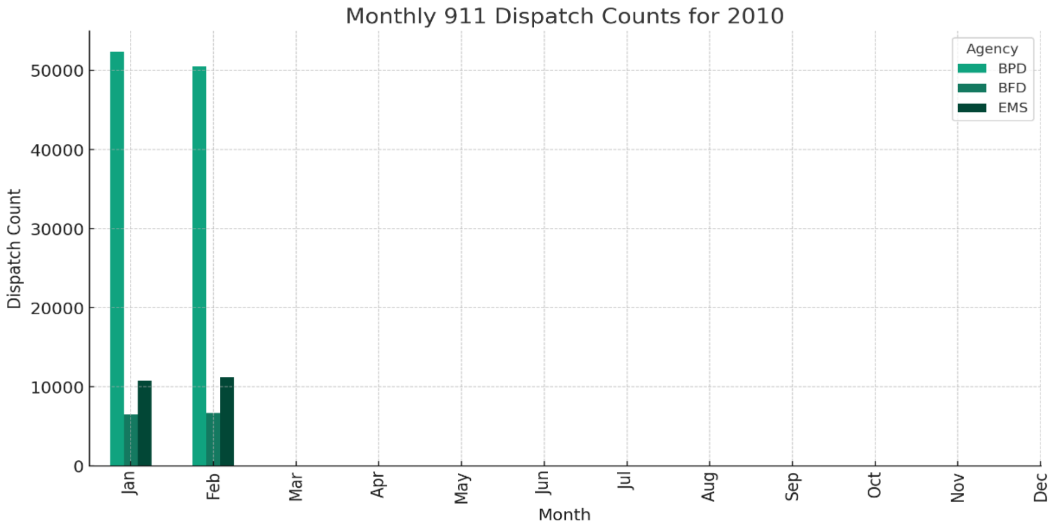
Monthly 911 Dispatch Counts for 2010
Dispatch Count
Month
Agency
BPD
BFD
EMS
0
10000
20000
30000
40000
50000
Jan
Feb
Mar
Apr
May
Jun
Jul
Aug
Sep
Oct
Nov
Dec

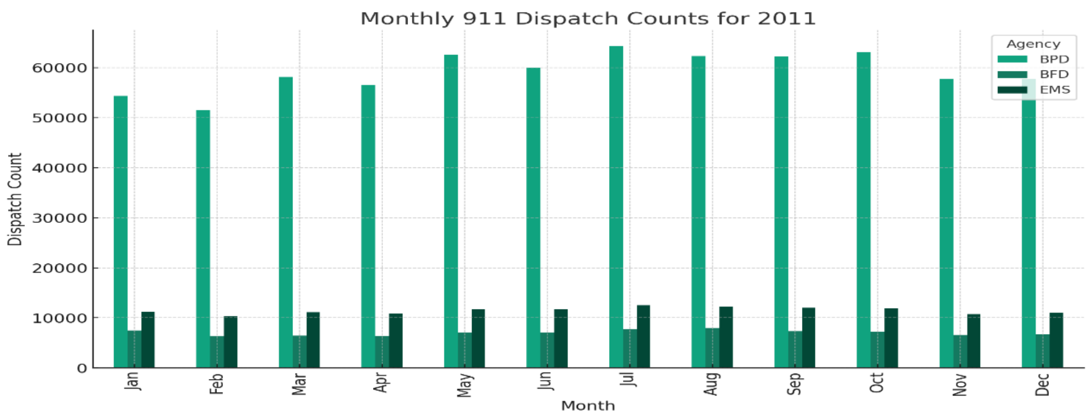
Monthly 911 Dispatch Counts for 2011
Dispatch Count
Month
Agency
BPD
BFD
EMS
0
10000
20000
30000
40000
50000
60000
Jan
Feb
Mar
Apr
May
Jun
Jul
Aug
Sep
Oct
Nov
Dec

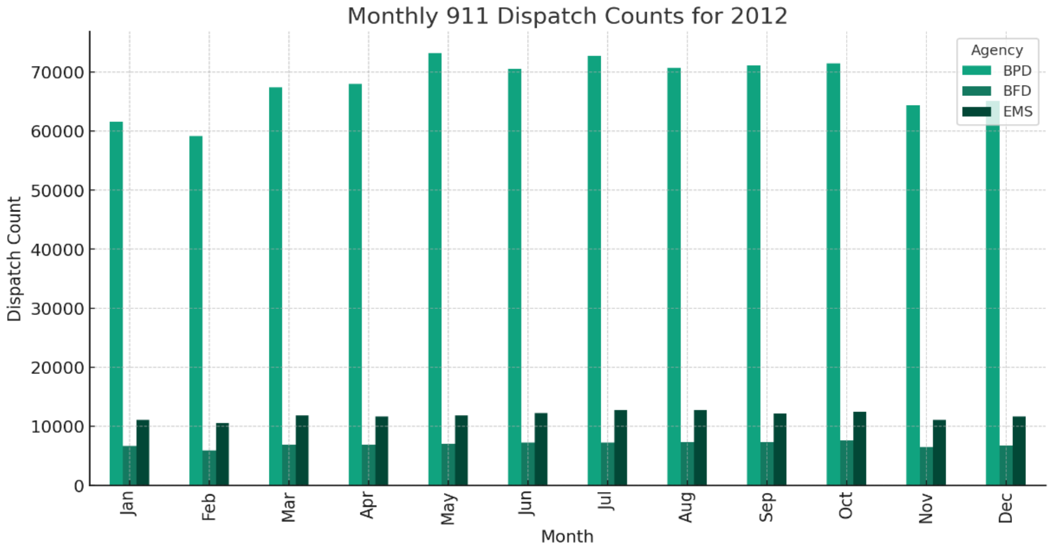
Monthly 911 Dispatch Counts for 2012
Agency
BPD
BFD
EMS
Dispatch Count
0
10000
20000
30000
40000
50000
60000
70000
Jan
Feb
Mar
Apr
May
Jun
Jul
Aug
Sep
Oct
Nov
Dec
Month

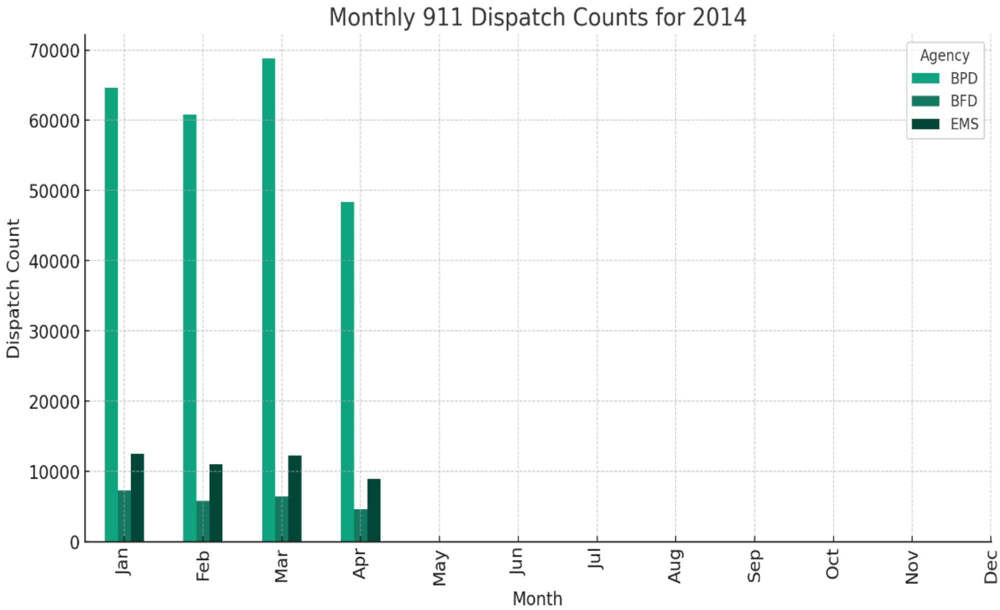
Monthly 911 Dispatch Counts for 2014
Agency
BPD
BFD
EMS
Dispatch Count
0
10000
20000
30000
40000
50000
60000
70000
Jan
Feb
Mar
Apr
May
Jun
Jul
Aug
Sep
Oct
Nov
Dec
Month

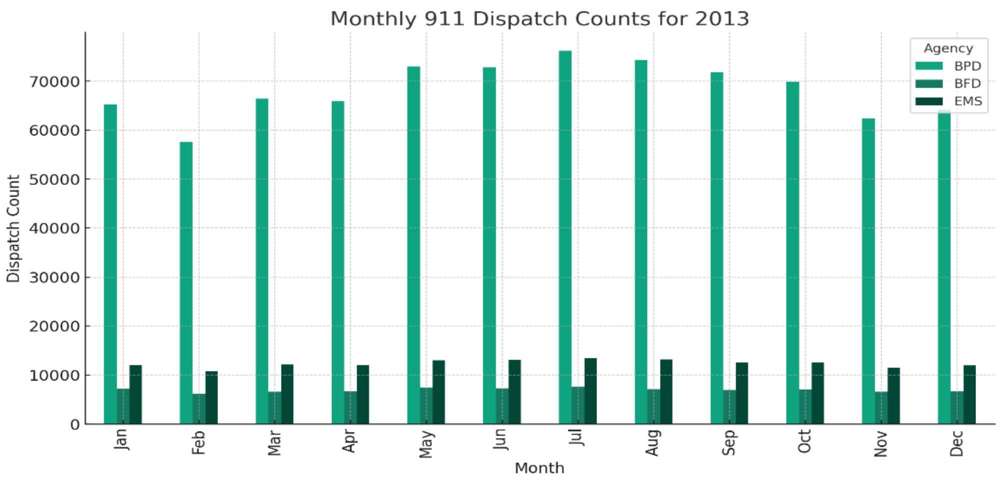
Monthly 911 Dispatch Counts for 2013
Agency
BPD
BFD
EMS
Dispatch Count
0
10000
20000
30000
40000
50000
60000
70000
Jan
Feb
Mar
Apr
May
Jun
Jul
Aug
Sep
Oct
Nov
Dec
Month

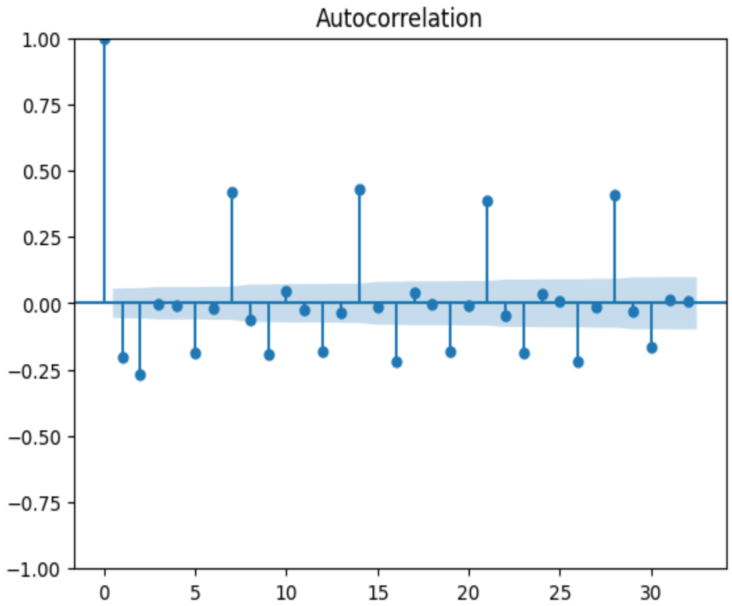

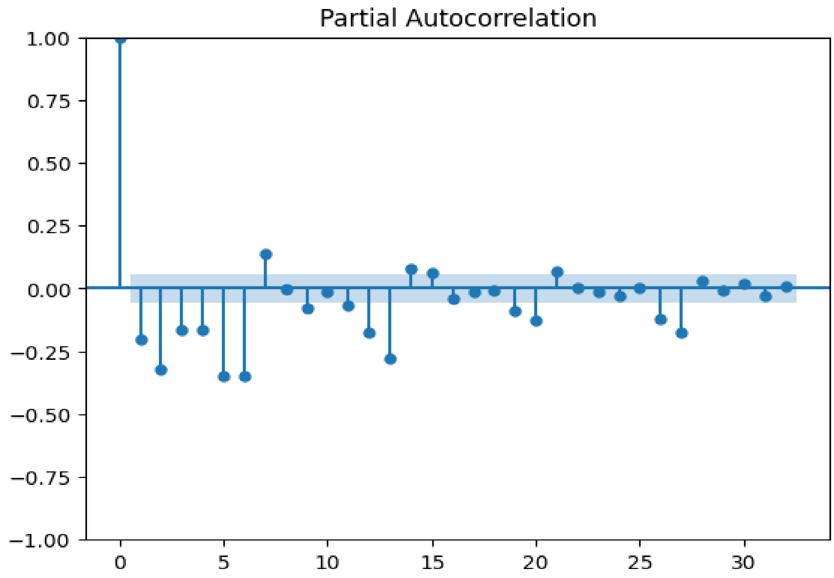

<u>Diagnostic Plots for Time Series Model Residuals:</u>

These diagnostic plots offer a comprehensive evaluation of the fitted time series model's residuals. The standardized residuals over time display randomness, indicating model adequacy. The histogram and KDE suggest a fair approximation to normality, though a slight deviation is observed, hinting at potential model improvements. The Q-Q plot largely confirms the residuals' normality, with most points aligning well with the theoretical quantiles. Finally, the correlogram demonstrates minimal autocorrelation in the residuals, affirming the model's ability to capture the data's temporal structure. Collectively, these diagnostics support the model's effectiveness while also highlighting areas for further refinement.

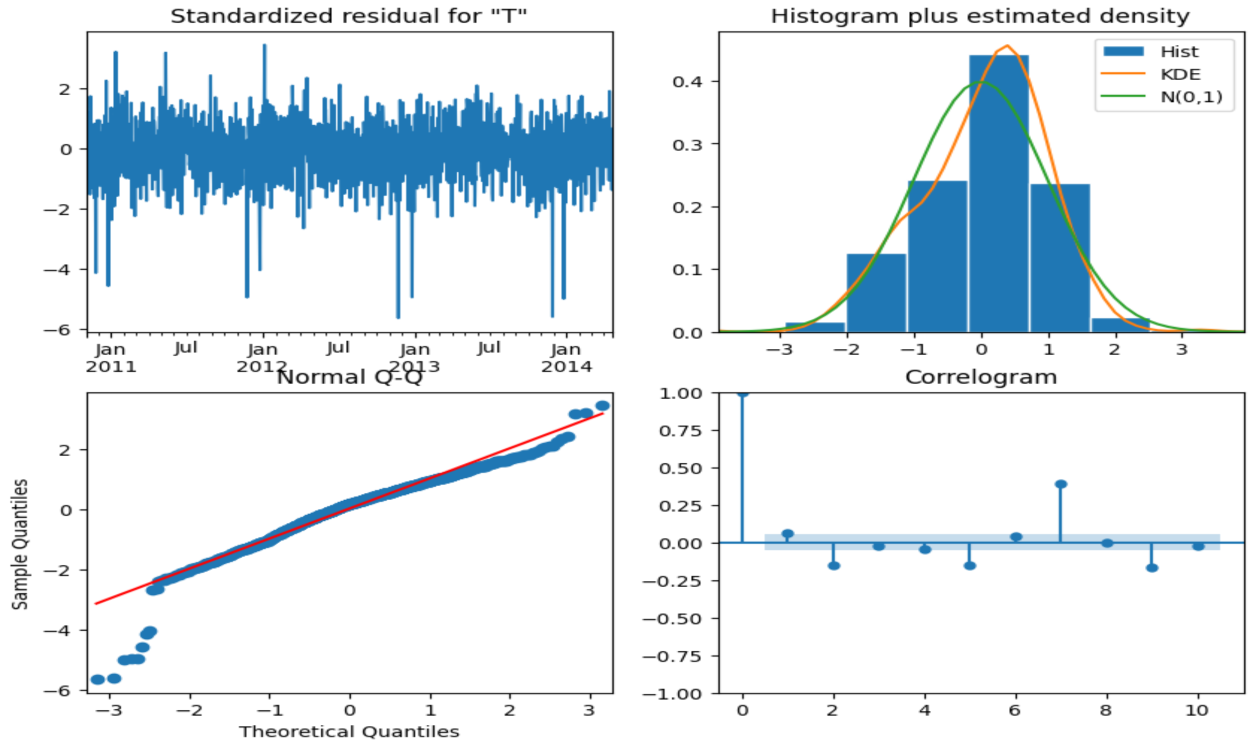

Stationarity Testing and Time Series Analysis:

Augmented Dickey-Fuller (ADF) Test Output:

ADF Test p-value: 0.135

The test indicates non-stationarity within the 911 dispatch count data (fail to reject the null hypothesis).

Kwiatkowski-Phillips-Schmidt-Shin (KPSS) Test Output:

KPSS Test p-value: 0.01

The KPSS test suggests non-stationarity (reject the null hypothesis), indicating a trend-stationary series.

Differencing Transformation:

After applying first differencing to the data, the ADF test yielded a significantly low p-value, suggesting that the transformed series is stationary.

ADF Test p-value 1.44×10^{-21}

This strongly indicates stationarity in the differenced series, confirming the effectiveness of the transformation.

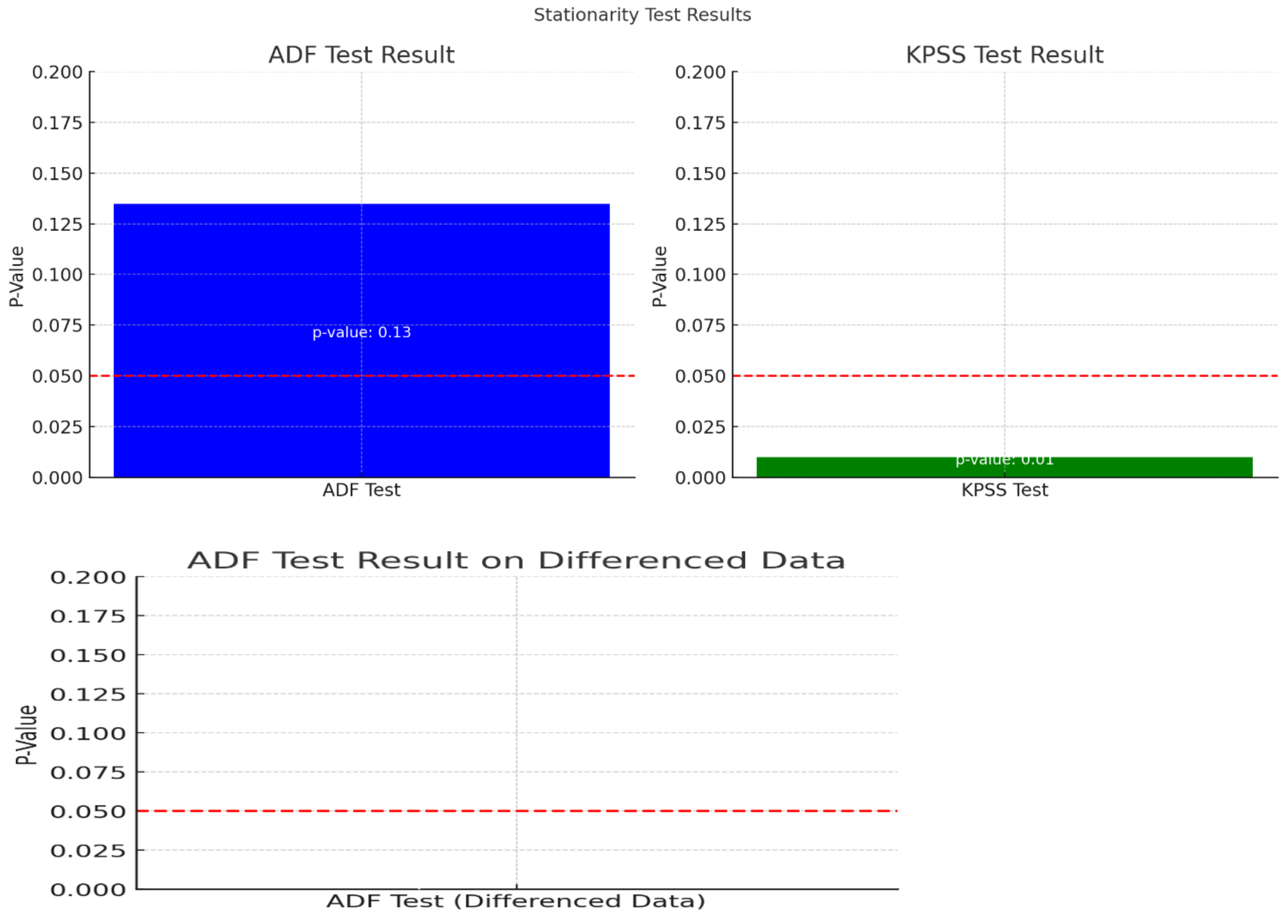

Evaluation of ARMA Model Fit to Differenced 911 Dispatch Data:

The graph presents the fitting of an ARMA (Autoregressive Moving Average) model to the differenced dispatch count data from a 911 call dispatch center. The actual data points are depicted in teal, while the model's fitted values are overlaid in red. This visualization enables us to compare the model's predictions against the observed data over time, from January 2011 through the end of 2013.

The plot indicates that the ARMA model, while capturing the general volatility in dispatch counts, may not completely grasp the extremities and certain movements within the series. The substantial overlap between actual and fitted values suggests that the model is responsive to the series fluctuations, yet the presence of peaks and troughs outside the fitted line suggests room for improvement. This could involve refining the model parameters, incorporating additional data, or considering alternative models that might account for potential seasonal patterns or exogenous variables affecting dispatch counts.

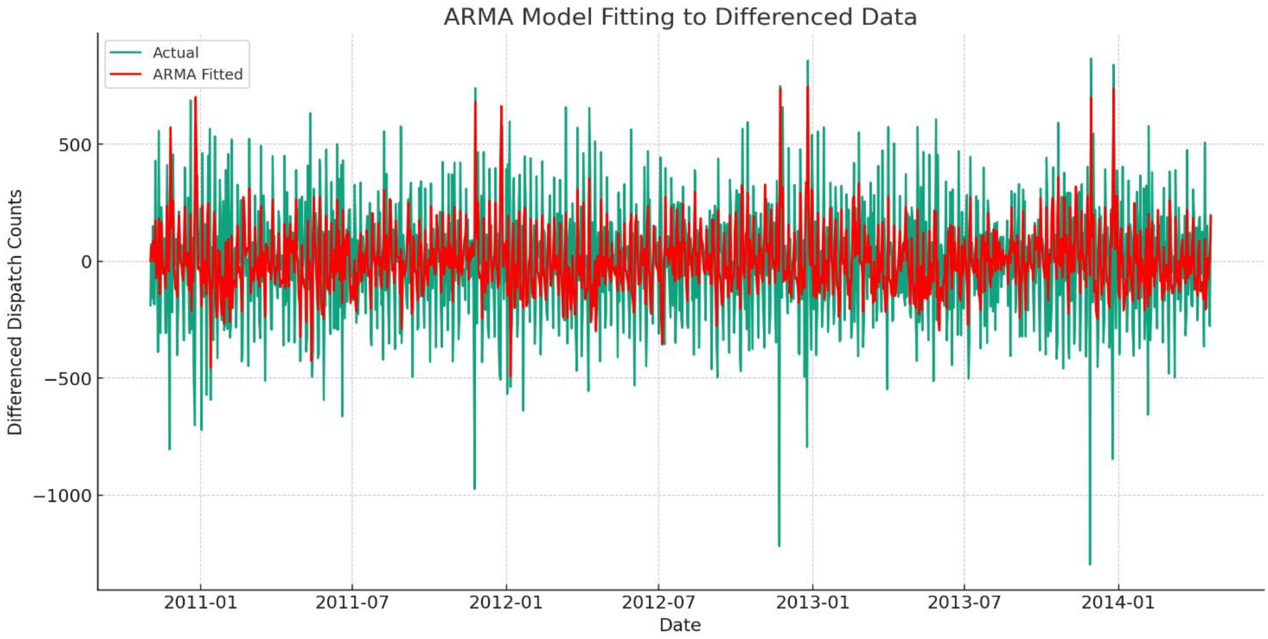

Evaluation of SARIMA Model Fit to Differenced 911 Dispatch Data:

This below graph showcases the application of a SARIMA (Seasonal AutoRgressive Integrated Moving Average) model to forecast 911 dispatch counts. The training data (in blue) was used to develop the model, capturing the historical patterns in dispatch activity. The actual observed data (in orange) from the test set represents the true values we aim to predict. The forecasted data (in green), along with the confidence interval (shaded area), indicates the model's predictions and its uncertainty. The plot reveals the model's ability to track the overall level of dispatch counts, though the variability of the actual data suggests potential for refinement. The confidence interval widens over time, reflecting increasing uncertainty in the forecasted values the further out we predict. The model's efficacy is quantified by an RMSE of approximately 366.23, which measures the average magnitude of the forecast errors.

This visualization and the accompanying metrics allow for an evaluation of the model's performance and provide a basis for further model tuning or the exploration of additional predictive factors that might improve forecast accuracy. The SARIMA model's predictions are crucial for operational planning and resource allocation within emergency dispatch management.

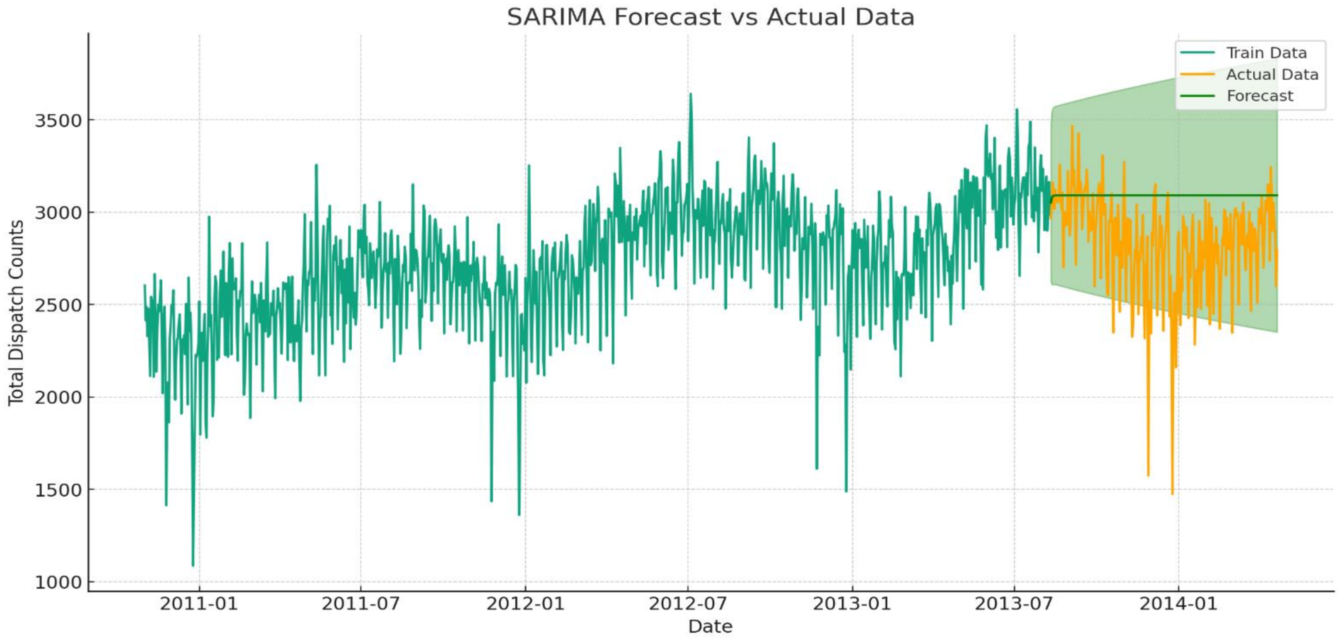

Appendix: C Code

```
import matplotlib.pyplot as plt
import seaborn as sns
import pandas as pd
from statsmodels.tsa.statespace.sarimax import SARIMAX
from google.colab import files
import io
uploaded = files.upload()
filename = next(iter(uploaded))

df = pd.read_csv(io.BytesIO(uploaded[911-daily-dispatch-count-by-agency.csv]))

print(df.columns)

df['Date'] = pd.to_datetime(df['Date'])
df.set_index('Date', inplace=True)

column_name = 'BPD'
split_point = int(len(df) * 0.8)
train, test = df.iloc[:split_point], df.iloc[split_point:]

sarima_model = SARIMAX(train[column_name], order=(1, 1, 1),
seasonal_order=(1, 1, 1, 12))
sarima_result = sarima_model.fit(disp=False)

model_params = sarima_result.params
model_conf_int = sarima_result.conf_int()
```

```python
model_pvalues = sarima_result.pvalues
model_zscores = sarima_result.zvalues
model_stderr = sarima_result.bse

model_results = {
    "Parameter": model_params.index,
    "Coefficient": model_params.values,
    "Standard Error": model_stderr.values,
    "Z-Score": model_zscores.values,
    "P-Value": model_pvalues.values,
    "Confidence Interval (Lower)": model_conf_int.iloc[:, 0].values,
    "Confidence Interval (Upper)": model_conf_int.iloc[:, 1].values
}

df_results = pd.DataFrame(model_results)

# Plotting the coefficients with confidence intervals
plt.figure(figsize=(10, 6))
sns.barplot(x='Coefficient', y='Parameter', data=df_results,
capsize=.2)
plt.errorbar(x=df_results['Coefficient'], y=df_results.index,
             xerr=[df_results['Coefficient'] - df_results['Confidence
Interval (Lower)'],
                   df_results['Confidence Interval (Upper)'] -
df_results['Coefficient']],
             fmt='none', c='black', capsize=5)
plt.title('SARIMAX Model Coefficients with Confidence Intervals')
plt.xlabel('Coefficient Value')
plt.ylabel('Parameter')
plt.grid(True)
plt.show()

# Plotting Z-Scores
plt.figure(figsize=(10, 6))
sns.barplot(x='Z-Score', y='Parameter', data=df_results)
plt.title('Z-Scores of SARIMAX Model Parameters')
plt.xlabel('Z-Score Value')
plt.ylabel('Parameter')
plt.grid(True)
plt.show()
```

Contributions:

Manoj Sankuru: Contributed to the Issues, Discoveries, Methods, Coding, and Outcome sections. Worked on graphs, the findings and utilized different regression analysis techniques and evaluations to interpret the data.

Medipalli Anji Reddy: Engaged with the Issues, Conversations, Methods, and Outcomes sections. Independently graphed data interpretations using the techniques described in the document.

Rama Satya Sai Prasad Appari: Focused on the issues and crafting the code for analytical models. He was also responsible for generating the graphics based on these models.

Suram Karthik Reddy: Delved into the preliminary investigations, refined the data, and explored various modelling options. Tested the efficiency of different fits and evaluated non-linear models to interpret the relationships within predictors.