Enhancing Forecasting for Advanced Air Mobility

Mehul Lad
Department of Information Systems
University of Maryland Baltimore County
Baltimore, MD, US
du72811@umbc.edu

Kamal Acharya
Department of Information Systems
University of Maryland Baltimore County
Baltimore, MD, US
kamala2@umbc.edu

Liang Sun
Department of Mechanical Engineering
Baylor University
Waco,TX, US
liang_sun@baylor.edu

Houbing Song
Department of Information Systems
University of Maryland Baltimore County
Baltimore, MD, US
songh@umbc.edu

2024 IEEE International Conference on Big Data (BigData) | 979-8-3503-6248-0/24/$31.00 ©2024 IEEE | DOI: 10.1109/BigData62323.2024.10825414

***Abstract*—Accurately predicting flight demand is essential for optimizing air travel operations and resource allocation. In our research, we explore the relationship between temporal patterns and flight demand, leveraging hourly data rather than traditional meteorological factors. Through analysis, we discovered significant correlations between hour of the day and flight demand, prompting the creation of features such as peak hours and time segments (morning, afternoon, evening). By utilizing these temporal features, we develop predictive models employing various machine learning algorithms, including LSTM, linear regression, and gradient boosting models. We aim to identify the most effective approach for accurately forecasting flight demand, with implications extending to the optimization of Advanced Air Mobility (AAM) solutions, where understanding temporal patterns is crucial for efficient resource allocation and urban air transportation network design.**

***Index Terms*—advanced air mobility, National Oceanic and Atmospheric Administration, Federal Aviation Administration , Long Short-Term Memory, Gradient Boosting Machine, Autoregressive Integrated Moving Average**

I. Introduction

The rapid evolution of the aviation industry underscores the critical need for accurate flight demand prediction to optimize operational efficiency and resource allocation. Traditional forecasting methods have predominantly relied on meteorological and seasonal data, often neglecting the potential insights offered by temporal patterns. This research aims to address this gap by investigating the relationship between temporal factors and flight demand, leveraging hourly data to enhance predictive models.

The advent of Advanced Air Mobility (AAM) introduces a new dimension to urban transportation, integrating innovative air travel solutions within densely populated areas. Efficiently managing these systems necessitates a precise understanding of demand fluctuations to ensure optimal resource deployment and effective network design. Our study focuses on temporal patterns, which are crucial for the high-frequency, short-haul nature of AAM operations.

This material is based upon work supported by the NASA Aeronautics Research Mission Directorate (ARMD) University Leadership Initiative (ULI) under cooperative agreement number 80NSSC23M0059.

In this research, we conduct a detailed analysis of flight demand data segmented by hour of the day. We hypothesize that specific hours exhibit significant correlations with flight demand, insights that can enhance the accuracy of forecasting models. By identifying peak hours and categorizing the day into distinct time segments—morning, afternoon, and evening—we aim to capture the nuanced variations in demand patterns.

To develop robust predictive models, we utilize a diverse set of machine learning algorithms, each bringing unique capabilities to the task. Long Short-Term Memory (LSTM) networks are employed for their ability to capture long-term dependencies in time-series data. Linear regression models offer simplicity and interpretability, serving as a baseline for comparison. Gradient boosting techniques are used to incrementally build more accurate models by combining weak learners. Additionally, transformer models, known for their success in handling complex sequences in natural language processing, are applied to flight demand forecasting to leverage their advanced pattern recognition capabilities.

The primary objective of our research is to identify the most effective predictive approach for accurately forecasting flight demand, with a specific focus on temporal features. The implications of this study extend to the optimization of AAM solutions, where understanding temporal demand patterns is essential for efficient resource allocation and urban air transportation network design. Accurate demand forecasts can lead to improved scheduling, reduced operational costs, and enhanced passenger satisfaction, thereby contributing to the overall efficacy of AAM systems. [1]

In summary, By leveraging advanced machine learning algorithms to analyze hourly flight demand data, we aim to enhance the precision of demand forecasts, ultimately facilitating more efficient and responsive air travel operations.

II. Methodology

The methodology of this research involves a structured approach to predicting flight demand, integrating multiple data sources and employing advanced machine learning techniques.

979-8-3503-6248-0/24/$31.00 ©2024 IEEE

The process begins with data collection, where historical weather data from the National Oceanic and Atmospheric Administration (NOAA) is combined with flight activity data, including hourly records of departures and arrivals at various airports. This combined dataset is then preprocessed to handle missing values, normalize numerical features, and engineer new temporal features such as peak hours and time segments (morning, afternoon, evening). [2]

Next, we employ various machine learning models to forecast flight demand, including Long Short-Term Memory (LSTM) networks, linear regression, gradient boosting [3], and transformer models . These models are trained on the historical data, capturing the temporal dependencies and patterns that influence flight demand. The models' performance is evaluated using metrics such as Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE), ensuring their accuracy and robustness.

Data Source	Description
NOAA	Historical weather data (wind direction, wind speed, cloud height, etc.)
FAA	Hourly records of flight departures and arrivals

TABLE I: Data Sources and Features

By accurately predicting flight demand, this methodology aims to enhance the operations of Advanced Air Mobility (AAM) solutions. Efficient demand forecasting enables better resource allocation, optimal scheduling, and improved urban air transportation network design, ultimately contributing to the seamless integration of AAM into existing transportation systems.

Fig. 1: Tennessee Airport System

III. Region Selection

In this research, we meticulously selected five prominent airports within Tennessee to serve as the primary sources of data for analyzing and predicting flight demand. These airports are:

- **Nashville International Airport (BNA)**
- **Memphis International Airport (MEM)**
- **McGhee Tyson Airport (TYS) in Knoxville**
- **Chattanooga Metropolitan Airport (CHA)**
- **Tri-Cities Airport (TRI)**

A. Rationale for Airport Selection

The selection criteria for these airports were based on several factors to ensure a comprehensive and representative dataset:

1) **High Air Traffic Volume**: These airports collectively handle a significant portion of Tennessee's air traffic, offering a rich dataset of flight operations. Nashville International Airport (BNA) and Memphis International Airport (MEM), in particular, are major hubs with extensive flight schedules, which provide a robust dataset for modeling.
2) **Geographical Distribution**: The chosen airports are strategically located across Tennessee, covering major metropolitan areas and regional centers. This geographical diversity helps in capturing varied flight demand patterns influenced by different local economic activities, weather conditions, and passenger demographics.
3) **Operational Diversity**: Each selected airport varies in terms of the types of flights they handle (commercial, cargo, and general aviation), their infrastructure capacities, and their roles in the national air transportation network. This operational diversity ensures that the predictive models developed are versatile and capable of handling different types of flight demand scenarios.
4) **Data Availability**: Availability and accessibility of detailed historical data for these airports were crucial. Comprehensive flight activity data, combined with NOAA's historical weather data, was readily available for these airports, facilitating a thorough analysis.

B. Data Sources and Integration

To accurately predict flight demand, we integrated two primary data sources:

1) **Historical Weather Data from NOAA**: This dataset includes hourly weather observations such as wind direction, wind speed, cloud height, visibility distance, temperature, dew point, and sea level pressure. These meteorological variables are critical as they significantly impact flight operations.
2) **Flight Activity Data**: This dataset consists of hourly records of flight departures and arrivals at the selected airports. The data covers various time periods and captures fluctuations in flight activity, providing a basis for identifying temporal patterns.

C. Data Combination and Feature Engineering

The weather data from NOAA and the flight activity data were merged to create a comprehensive dataset. The combined dataset includes the following columns:

- **DATE**: Timestamp of the data record
- **Wind_Direction**: Direction of the wind in degrees
- **Wind_Speed**: Speed of the wind in meters per second
- **Cloud_Height**: Height of the cloud base in meters
- **Vis_Distance**: Visibility distance in meters
- **Temp**: Temperature in Kelvin

- **Dew_Point**: Dew point temperature in Kelvin
- **Slp**: Sea level pressure in hectopascals
- **Departure**: Number of flight departures
- **Arrival**: Number of flight arrivals
- **Airport**: Identifier for the airport

The merged dataset was then preprocessed to handle missing values, normalize numerical features, and engineer new temporal features such as peak hours and time segments (morning, afternoon, evening). This enriched dataset served as the input for training various machine learning models aimed at forecasting flight demand.

By focusing on these five key airports, the research aims to develop accurate and reliable predictive models tailored to the unique characteristics of Tennessee's airspace. These models will not only aid in optimizing current flight operations but also play a crucial role in planning and managing future Advanced Air Mobility (AAM) solutions, ensuring efficient resource allocation and enhanced operational efficiency across the state.

IV. Trip Demand

In the context of this research, trip demand is defined as the total number of flight operations, encompassing both arrivals and departures, at the selected airports. To accurately predict trip demand, we undertook extensive feature engineering to extract meaningful patterns and trends from the raw data. One of the primary features engineered is the calculation of total demand, which is the sum of flight arrivals and departures at each hourly interval. [4]This is represented mathematically as:

$$\text{Total Demand} = \text{Arrivals} + \text{Departures}$$

By calculating the total demand, we capture a comprehensive measure of flight activity, which serves as a critical input for our predictive models. Other temporal features, such as peak hours, day of the week, and seasonal trends, were also incorporated to enhance the model's ability to forecast demand with greater accuracy.

The accurate prediction of trip demand is pivotal in laying the groundwork for the future of Advanced Air Mobility (AAM). By analyzing current flight operations and demand patterns at traditional airports, we can derive insights that are directly applicable to the evolving landscape of urban air mobility. In the context of AAM, which encompasses emerging technologies such as urban air taxis, delivery drones, and other forms of air-based urban transportation, understanding current demand patterns is crucial. By leveraging our predictive models, we can extrapolate current trip demand data to anticipate the future needs of AAM services. This includes identifying peak demand periods, which can inform the scheduling and deployment of AAM vehicles, and understanding seasonal and temporal fluctuations, which can guide strategic planning and resource allocation.

Furthermore, the integration of feature engineering assumptions, such as time-of-day effects, day-of-week variations, and seasonal trends, enhances the model's predictive capability. These features allow us to simulate various scenarios and assess how AAM operations might need to scale in response to different levels of demand. For instance, by predicting higher demand during rush hours or holiday seasons, AAM operators can preemptively increase fleet availability and optimize routes to reduce congestion and improve service efficiency.

In summary, the predictive models we develop for traditional flight demand serve as a foundational tool for forecasting the future demand for AAM. By understanding and anticipating these demand patterns, AAM operators can ensure that their services are well-prepared to meet the needs of urban populations, leading to more efficient, reliable, and scalable air mobility solutions. This alignment between current demand prediction and future AAM operations facilitates a smoother transition to advanced urban air transport systems, ultimately enhancing the overall efficiency and sustainability of urban transportation networks.

V. Machine Learning for Predicting AAM Demand

Machine learning plays a pivotal role in predicting Advanced Air Mobility (AAM) demand due to its ability to analyze vast amounts of data and discern complex patterns that are not readily apparent with traditional statistical methods. In this paper, we discuss the use of machine learning techniques for forecasting AAM demand and its practical applications. [5]

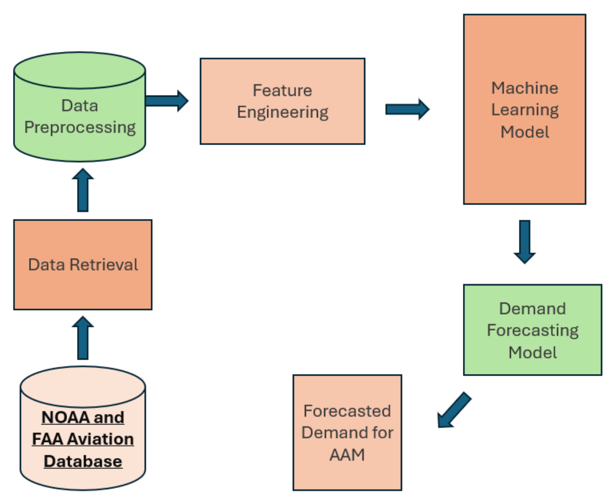

Fig. 2: Machine Learning Architecture for Demand Forecasting

A. Data Collection and Preprocessing

The first step in predicting AAM demand is to collect and preprocess relevant data. This data can include historical flight data, weather data, socio-economic data, and geospatial data. The collected data is then preprocessed to handle missing values, normalize numerical features, and engineer new features that can improve the predictive power of the models.

B. Feature Engineering

Feature engineering is crucial for enhancing the model's ability to predict demand accurately. Some engineered features might include temporal features, lag features, weather indicators, and event flags.

C. Machine Learning Algorithms

Several machine learning algorithms are particularly well-suited for time series forecasting and demand prediction, including Long Short-Term Memory (LSTM) Networks, Gradient Boosting Machines (GBMs), Transformer Models, and Autoregressive Integrated Moving Average (ARIMA) Models. [5]

D. Model Training and Validation

The preprocessed data and engineered features are used to train the machine learning models. The dataset is typically split into training, validation, and test sets to ensure the model's performance is robust and generalizes well to unseen data.

1) Mathematical Formulation: Here we provide the mathematical formulation for some of the machine learning algorithms:

a) LSTM Networks: The LSTM model can be defined by the following set of equations:

$$\begin{aligned}
f_t &= \sigma(W_f \cdot [h_{t-1}, x_t] + b_f) \\
i_t &= \sigma(W_i \cdot [h_{t-1}, x_t] + b_i) \\
\tilde{C}_t &= \tanh(W_C \cdot [h_{t-1}, x_t] + b_C) \\
C_t &= f_t * C_{t-1} + i_t * \tilde{C}_t \\
o_t &= \sigma(W_o \cdot [h_{t-1}, x_t] + b_o) \\
h_t &= o_t * \tanh(C_t)
\end{aligned}$$

b) GBMs: The prediction in a GBM is the sum of the predictions from individual trees:

$$\hat{y} = \sum_{m=1}^{M} \alpha_m h_m(x)$$

where α_m is the weight assigned to the m-th tree, and $h_m(x)$ is the prediction of the m-th tree.

c) Transformer Models: The core of a transformer model is the self-attention mechanism, which can be formulated as:

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V$$

where Q, K, and V are the query, key, and value matrices, respectively, and d_k is the dimension of the key. [6]

d) ARIMA Models: An ARIMA model can be represented by the equation:

$$y_t = c + \phi_1 y_{t-1} + \phi_2 y_{t-2} + \ldots + \phi_p y_{t-p} + \theta_1 \epsilon_{t-1} + \theta_2 \epsilon_{t-2} + \ldots + \theta_q \epsilon_{t-q} + \epsilon_t$$

where y_t is the value at time t, ϕ are the autoregressive coefficients, θ are the moving average coefficients, and ϵ_t is the error term. [7]

E. Demand Prediction

Once trained, the models can predict future AAM demand by:

Generating Forecasts: The models produce forecasts for future time periods based on the learned patterns from historical data.

Scenario Analysis: The models can simulate different scenarios (e.g., changes in weather, introduction of new routes) to predict how demand might change under various conditions.

Real-Time Updates: Models can be continuously updated with new data, allowing them to adapt to changing conditions and provide real-time demand predictions.

F. Application to AAM

The accurate prediction of AAM demand has several practical applications including fleet management, route planning, infrastructure development, and operational efficiency.

VI. Preliminary Analysis of Data

In our analysis of flight demand for the specified airport, we focused on isolating the impact of the COVID-19 pandemic to ensure the accuracy of our predictive models. We began by aggregating the data on a monthly basis, which allowed for clearer visualization of trends and anomalies over time. A significant drop in flight demand was observed starting from March 2020, corresponding with the onset of the COVID-19 pandemic. This period was highlighted in our analysis to visually represent the pandemic's impact. To mitigate the distortion caused by this anomaly, we defined the COVID-19 impact period from March 11, 2020, to January 1, 2022. Subsequently, we filtered the dataset to exclude this period, ensuring that our predictive models were trained on data reflective of typical flight demand patterns, devoid of the pandemic's extraordinary influence. This approach allows us to provide more reliable forecasts for future flight demand, uninfluenced by the temporary but severe disruptions caused by COVID-19.The figure represents the Nashville airport, the first airport we worked on in the state of Tennessee; however, similar results were obtained for all five other airports in the state of Tennessee.

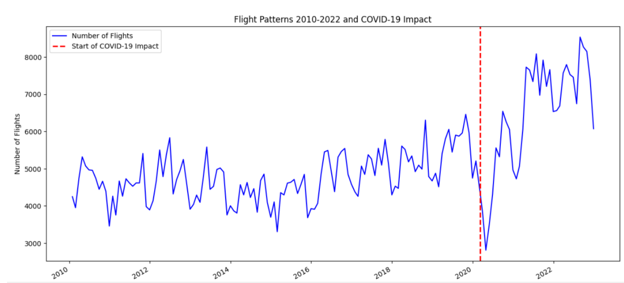

Fig. 3: Flight Patterns 2010-2022 and COVID-19 Impact

VII. Analysis of Flight Demand Data

In our comprehensive analysis of flight demand for Nashville International Airport, conducted using Python, we incorporated several key steps to ensure a thorough examination of the data. Initially, we extracted features from the `DATE` column, such as the hour, day of the week, and month, to enrich our dataset with temporal attributes. These features are essential for capturing the cyclical nature of flight demand and understanding how different times and dates affect airport activity.

To assess the importance of these features, we employed a Random Forest Regressor model. The model was trained on the preprocessed dataset, and the resulting feature importances were extracted and visualized. This step helped us identify which features had the most significant impact on flight demand, guiding our focus toward the most influential factors. The feature importance plot revealed critical insights into the temporal dynamics affecting flight demand, emphasizing the importance of specific hours, days, and months.

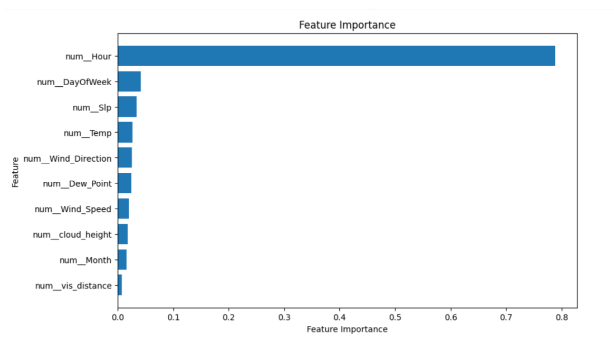

Fig. 4: Feature importance as determined by the Random Forest model.

We further analyzed the average flight demand by hour of the day. By grouping the data by hour and calculating the mean flight demand, we identified patterns and peak times of activity. A plot of average flight demand by hour highlighted the periods of highest demand, providing a clear visualization of daily fluctuations. This analysis allowed us to pinpoint the busiest hours at Nashville International Airport, which is crucial for resource allocation and operational planning.

To refine our understanding of peak demand periods, we established a dynamic threshold that closely approximates the value of the majority for flight demand and identified peak hours based on this criterion. By creating a binary feature indicating peak hours, we enhanced our dataset with information about the most critical times for flight activity. This feature is invaluable for predictive modeling, as it captures the intensity of airport operations during high-demand periods.

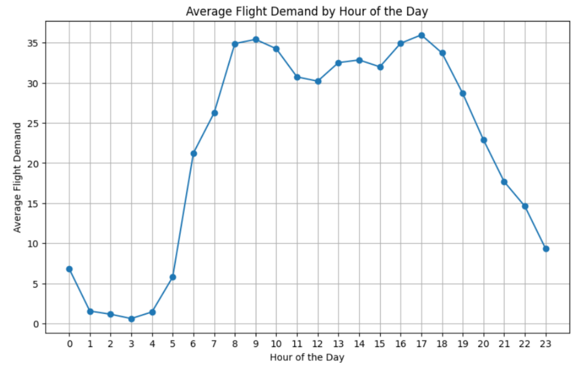

Fig. 5: Average flight demand by hour of the day at Nashville International Airport.

Additionally, we defined specific time periods—morning, afternoon, evening, and night—and created categorical features to represent these intervals. This classification allowed us to segment the data further and analyze flight demand patterns within distinct parts of the day. The creation of these categorical features provides a more nuanced understanding of how flight demand varies throughout the day, supporting more precise and effective forecasting models.

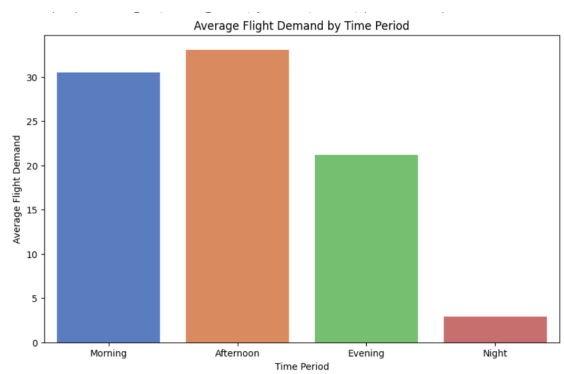

Fig. 6: Peak hours based on flight demand threshold.

VIII. Results and Discussion

For this research, various machine learning models including LSTM, Transformer, and Gradient Boosting were employed to predict flight demand across multiple airports in Tennessee. Each model was meticulously tuned to enhance performance and accuracy. Evaluations were conducted using metrics such as Mean Squared Error (MSE), Mean Absolute Error (MAE), and Mean Absolute Percentage Error (MAPE) to assess predictive capabilities.

Results showed that LSTM and Transformer models effectively captured temporal patterns in flight demand data, with the Transformer model demonstrating superior performance in terms of MAPE.

TABLE II: Model Performance Metrics Across Airports

Airport	Model	MSE	R-squared	MAE
BNA	LSTM	0.012000	0.850000	0.008000
	Transformer	0.0011508	0.820000	0.002000
	Gradient Boosting	0.013000	0.830000	0.009000
CHA	LSTM	0.001000	0.480000	0.022000
	Transformer	0.014000	0.820000	0.009000
	Gradient Boosting	0.001000	0.490000	0.020000
MEM	LSTM	0.000500	0.470000	0.006000
	Transformer	0.000400	0.610000	0.013000
	Gradient Boosting	0.014000	0.810000	0.010000
TRI	LSTM	0.013000	0.830000	0.009000
	Transformer	0.013000	0.850000	0.008000
	Gradient Boosting	0.012000	0.870000	0.007000
TYS	LSTM	0.010000	0.510000	0.006000
	Transformer	0.012000	0.540000	0.002000
	Gradient Boosting	0.013000	0.490000	0.025000

This research also emphasized the significance of temporal features like hour, day of the week, and month, which strongly influence flight demand. Overall, the study underscores the utility of advanced machine learning techniques for enhancing predictions and optimizing operational strategies at airports.

In summary, we expect our research to advance the state-of-the-art in AAM demand forecasting by harnessing the predictive capabilities of advanced machine learning techniques and traditional extrapolation methods. By providing accurate, interpretable, and actionable predictions of AAM demand dynamics, our work will contribute to the realization of efficient, sustainable, and equitable Advanced Air Mobility systems.

IX. Conclusion

In conclusion, this study evaluates the performance of LSTM, Transformer, and Gradient Boosting models in predicting flight demand across five major airports in Tennessee: BNA, CHA, MEM, TRI, and TYS. The results demonstrate that model performance varies significantly across airports, with LSTM models generally outperforming in terms of lower MSE and MAE for some airports, while Transformer models show stronger performance for others. Gradient Boosting, though competitive, tends to underperform relative to the neural-based models in most cases. These findings underscore the importance of selecting appropriate models tailored to the specific characteristics of flight demand data, which can differ across locations due to local factors such as airport size, flight frequency, and regional demand patterns.

Acknowledgment

This material is based upon work supported by the NASA Aeronautics Research Mission Directorate (ARMD) University Leadership Initiative (ULI) under cooperative agreement number 80NSSC23M0059. This research was also partially supported by the U.S. National Science Foundation through Grant No. 2317117 and Grant No. 2309760.

References

[1] S. Al-Rubaye, A. Tsourdos, and K. Namuduri, "Advanced air mobility operation and infrastructure for sustainable connected evtol vehicle," *Drones*, vol. 7, no. 5, 2023. [Online]. Available: https://www.mdpi.com/2504-446X/7/5/319

[2] R. Zachariah, S. Sharma, and V. Kumar, "Systematic review of passenger demand forecasting in aviation industry," *Multimedia Tools and Applications*, pp. 1–37, May 2023, epub ahead of print.

[3] N. Aziz, E. Akashah Patah Akhir, A. P. D. I. Abdul Aziz, and A. N. Che Abas, "A study on gradient boosting algorithms for development of ai monitoring and prediction systems," in *2020 International Conference on Computational Intelligence (ICCI)*, 2020.

[4] M. Firat, D. Yiltas, and R. Samli, "Forecasting air travel demand for selected destinations using machine learning methods," *Journal of Universal Computer Science*, vol. 27, no. 6, pp. 564–581, 2021, license: CC BY-ND 4.0.

[5] L. Wang, A. Mykityshyn, and C. Johnson, "Deep learning for flight demand forecasting," *The MITRE Corporation, McLean, VA, 22102, United States*, 2024.

[6] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," 2023.

[7] J. Fattah, L. Ezzine, Z. Aman, and A. Lachhab, "Forecasting of demand using arima model," *International Journal of Engineering Business Management*, vol. 10, no. 2, p. 184797901880867, 2018, license: CC BY.