

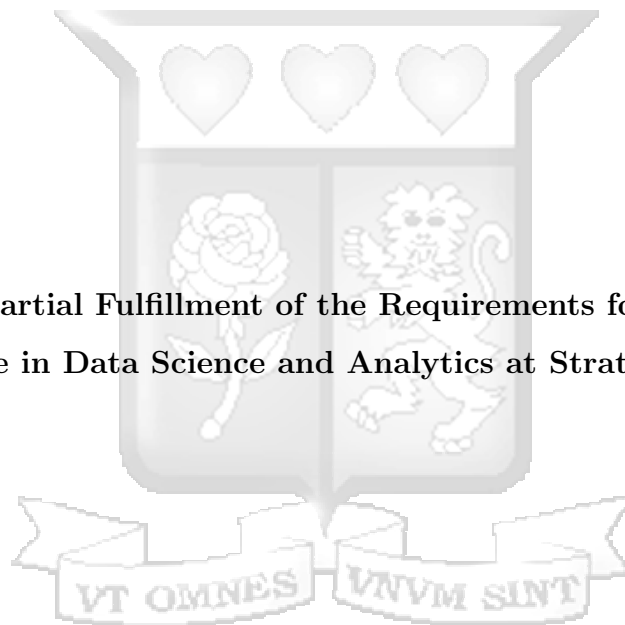
**Predictive Maintenance of Aircraft Engines Using Machine Learning: A
Case of the CFM56-7B26E Engine**

By

Cecilia Mainga Oketch

168951

**Submitted in Partial Fulfillment of the Requirements for the Degree of
Master of Science in Data Science and Analytics at Strathmore University**



**Institute of Mathematical Sciences
Strathmore University
Nairobi, Kenya**

June, 2025

This dissertation is available for library use on the understanding that it is copyright material and that no quotation from the dissertation may be published without proper acknowledgement.

Declaration and Approval

Declaration

I declare that this work has not been previously submitted and approved for the award of a degree by this or any other University. To the best of my knowledge and belief, the dissertation contains no material previously published or written by another person except where due reference is made in the dissertation itself.

©No part of this dissertation may be reproduced without the permission of the author and Strathmore University.

Student's Name: **Cecilia Mainga Oketch**

Sign: 

Date: 23/5/2025

Approval

The dissertation of **Cecilia Mainga Oketch** was reviewed and approved by the following:

Dr. Kennedy Senagi,
Lecturer, Institute of Mathematical Sciences,
Strathmore University.

Dr. Godfrey Madigu,
Dean, Institute of Mathematical Sciences,
Strathmore University.

Prof. Bernard Shibwabo,
Director of Graduate Studies,
Strathmore University.

Abstract

Understanding the Exhaust Gas Temperature margin is essential for predictive maintenance of aircraft engines. It addresses the challenges of unscheduled maintenance and equipment downtime. Since aircraft engines are expensive to purchase and maintain, airline operators strive to maximize the use of installed engines. This project used several machine learning models to predict the EGT Hot Day Margin. A real CFM56-7B26E engine dataset is used for analysis and prediction. The models were evaluated using metrics such as mean absolute error, mean square error and root mean square error. Results showed that the deep learning model performed better than all the other regression algorithms, and cross-validation against engine operational limits achieved an accuracy of 98%. SHAP analysis further identified the most relevant features influencing the model predictions. The key factors include EGT, indicated fan Speed, total air temperature, core speed, mach, and the Operational length of the engine, highlighting their significant impact on engine performance. The best-performing model model was then deployed on the Streamlit community cloud platform. This project demonstrates how machine learning can enhance the reliability of aircraft engine maintenance practices.

Keywords: Exhaust Gas Temperature(EGT), Predictive Maintenance, Aircraft Engine, Machine Learning.



Table of Contents

Declaration and Approval	ii
Abstract.	iii
List of Figures	vii
List of Tables.	viii
List of Abbreviations	ix
Acknowledgment	x
Chapter 1: Introduction	1
1.1 Background	1
1.2 Problem Statement	2
1.3 Research Aim	3
1.4 Research Objectives	3
1.5 Research Questions	3
1.6 Scope and Limitations of the Study	3
1.6.1 Scope	3
1.6.2 Limitations	4
1.7 Research Justification	4
Chapter 2: Literature Review	5
2.1 Overview of the Operation of Aircraft Engines	5
2.2 Related Works	7
2.3 Research Gap	10
Chapter 3: Methodology	12
3.1 Business Understanding	13
3.2 Data Understanding	14
3.3 Data Preparation	17
3.3.1 Handling Missing Values	17
3.3.2 Feature Engineering	18
3.3.3 Dealing with Outliers	18
3.4 Machine Learning Modeling	18
3.5 Performance Evaluation Criteria	20
3.5.1 Regression Models Performance Evaluation	20
3.5.2 Model vs. Operational Limits Comparison	21

3.5.3	Shapley Additive Explanations (SHAP) Analysis	22
3.6	Deployment	23
Chapter 4:	System Design and Architecture.	24
4.1	System Modeling	24
4.2	System Components	25
4.2.1	Database	25
4.2.2	Web portal	27
Chapter 5:	System Implementation and Testing	30
5.1	System Implementation	30
5.1.1	Database	30
5.1.2	Web portal	30
5.1.3	Email Notification	34
5.2	Testing	34
5.2.1	Functionality Testing	35
5.2.2	Usability Testing	35
5.2.3	Compatibility Testing	35
5.2.4	Security Testing	36
Chapter 6:	Discussion of Results	37
6.1	Explanatory Data Analysis Results	37
6.2	Machine Learning Results	39
6.2.1	Class Imbalance	40
6.2.2	Comparison with Industry Limits	42
6.3	SHAP Analysis Results	44
6.3.1	SHAP Summary Plot	44
6.3.2	SHAP Feature Importance plot	46
6.3.3	SHAP Waterfall Plot	47
6.4	Summary	54
Chapter 7:	Conclusions and and Future Work	55
7.1	Conclusion	55
7.2	Future Work	55
Bibliography		57

Appendices 62
Appendix A: Similarity Report 62
Appendix B: Ethical Clearance Confirmation 64



List of Figures

2.1	Forces of flight	5
2.2	Major components of a Jet Engine	6
2.3	Illustration of EGT margin	6
3.1	CRISP-DM framework	12
4.1	UML diagram	24
4.2	Entity Relationship Diagram (ERD) of the tables	26
4.3	Sitemap of the web portal	27
4.4	Overview page wireframe	28
4.5	Predictions page wireframe	28
4.6	Interpretation page wireframe	29
5.1	Overview page	31
5.2	Predictions page	32
5.3	Interpretation page	33
5.4	Email Notification	34
6.1	Boxplot of numerical columns before scaling	37
6.2	Boxplot after min-max scaling	38
6.3	Correlation heatmap	39
6.4	Training and Validation loss plot for the original deep learning model	40
6.5	Training and validation loss plot of the optimized deep learning model	41
6.6	Class distribution plot	41
6.7	Confusion matrix of original deep learning model	42
6.8	Confusion Matrix of the Optimized deep learning model	43
6.9	SHAP summary plot	45
6.10	SHAP feature importance plot	48
6.11	SHAP Waterfall plot	49
6.12	Random forest force plot	51
6.13	Decision tree force plot	51
6.14	KNN force plot	51
6.15	Linear regression force plot	52
6.16	Support vector regressor force plot	52
6.17	SHAP decision plots	53

List of Tables

3.1	A description of engine performance parameters	15
4.1	Database Tables	26
6.1	Model performance metrics	39
6.2	Classification Report of the Original Deep Learning Model	43
6.3	Classification report of the optimized deep learning model	43



List of Abbreviations

ACARS Aircraft Communications Addressing and Reporting System

CRISP-DM Cross-Industry Standard Process for Data Mining

EGT Exhaust Gas Temperature

EGTM Exhaust Gas Temperature Margin

ERD Entity Relationship Diagram

HPC High Pressure Compressor

HPT High Pressure Turbine

IATA International Air Transport Association

ICAO International Civil Aviation Organization

KNN K-Nearest Neighbors

LSTM Long-Short-Term Memory

LPC Low Pressure Compressor

LPT Low Pressure Turbine

MAE Mean Absolute Error

MRO Maintenance, Repair and Overhaul

NARX Nonlinear AutoRegressive with eXogenous inputs

NASA National Aeronautics and Space Administration

OEM Original Equipment Manufacturer

RMSE Root Mean Square Error

SHAP SHapley Additive Explanations

SVM Support Vector Machine

TAT Total Air Temperature

UML Unified Modeling Languages



Acknowledgments

I am grateful to God for His provision, grace, and for granting me good health throughout the study period.

I am grateful to my family for their unwavering support and cheering me on throughout the study period.

I am grateful to my supervisor, Dr. Kennedy Senagi for his availability, guidance, and support without which this dissertation would not have been a success.

Finally, I am grateful to my classmates for their support and encouragement.



Chapter 1: Introduction

1.1 Background

The world's first scheduled commercial flight took place on January 1st, 1914, from St. Petersburg to Tampa, Florida. The seaplane, a Benoist Type XIV, carried only two people: the pilot and one passenger [Reilly \(1996\)](#). In the present day, air transport plays a significant role in the global transport of people and cargo. According to International Civil Aviation Organization ([ICAO](#)), the total number of passengers on scheduled airline services reached 3.3 billion in 2022, a 42.8% increase from the previous year, while the number of departures increased to 29.8 million, a 23.1% increase from the previous year. Out of this, Africa constitutes 2.2% of world traffic, an increase of 83.3% from the previous year. Moreover, 56.5 million freight tonnes were transported, demonstrating the essential role of air transport in global commerce [International Civil Aviation Organization \(2022\)](#). In Kenya, between 2017 and 2018, the overall number of passengers handled increased by 16.8%, from 10.1 million to 11.8 million. There was 23.3% increase in commercial cargo volume, from 290.8 thousand tonnes in 2017 to 358.7 thousand tonnes in 2018 [Kenya Civil Aviation Authority \(2024\)](#).

The evolution of aircraft engines and design has directly influenced the expansion of global passenger and cargo transport. As the primary source of propulsion, engines provide the power necessary to drive aircraft forward, enabling advancements in speed, efficiency, and capacity. From small piston-engine aircraft to modern high-capacity, fuel-efficient jets, aviation continues to evolve to meet growing demands while reducing environmental impact [Sehra and Whitlow \(2004\)](#).

However, the financial load of maintaining such a growing fleet is significant. According to International Air Transport Association ([IATA](#)), which represents 290 airlines globally, the global fleet consisted of 32,070 aircrafts in 2022, with 78% of them in service. Maintenance, Repair and Overhaul ([MRO](#)) services accounted for \$76.8 billion in airline spending globally, which is 10.9% of overall operating costs. The continued performance and safety of the aircraft depend on these MRO services. However, they present a significant financial challenge for airlines, especially as they seek to optimize efficiency. As global passenger traffic is projected to double by 2040, the demand for more efficient and

cost-effective maintenance solutions also grows [Association \(2022\)](#).

In order to streamline these [MRO](#) services, modern aircrafts feature a variety of sensors that capture extensive systems health monitoring data [Lee and Mitici \(2023\)](#). In the case of a Boeing 737, a two-engine commercial aircraft, each engine generates 20 terabytes of information per flight hour [Badea et al. \(2018\)](#), while an Airbus A380 has approximately 25,000 sensors recording sensor data across their components in flight [Stanton et al. \(2023\)](#). This increase in the sensor technology onboard modern airliners unlocks many opportunities for using data-driven approaches to predict aircraft system failure.

1.2 Problem Statement

Buying and maintaining an aircraft engine is costly. For instance the CFM56-5 engine costs approximately \$4.5 million, with each maintenance shop visit costing approximately \$1 million [Cai et al. \(2016\)](#). Concurrently, engine maintenance takes up 30% of the total cost of the aircraft maintenance [Keivanpour and Kadi \(2015\)](#). An engine shop visit occurs for two main reasons; replacement or repair of life limited parts, or when the Exhaust Gas Temperature Margin (Exhaust Gas Temperature Margin ([EGTM](#))) reaches a critical value. The [EGTM](#) is a critical indicator of engine health, and when it drops below a threshold value, it signifies potential engine failure that would require a shop visit. A spare engine is required during a shop visit, hence airlines try to optimize installed engines to avoid several shop visits occurring simultaneously. It is possible to need just one spare engine if shop visits do not overlap [Cai et al. \(2016\)](#). With the development of predictive maintenance, [EGTM](#) can be monitored and tracked to aid in maintenance operations optimization.

In the context of a Kenyan airline, [EGTM](#) forecasting relies solely on the [EGT](#) measurements. This approach is inefficient because the [EGT](#) margin is the difference between the maximum allowable [EGT](#) by the manufacturer and the [EGT](#) measured in flight. This indicates the existing method uses a variation of the [EGT](#) margin to predict the same [EGT](#) which may not yield accurate results. Additionally, [EGT](#) is a collective result of the operation of the engine, from the intake to the exhaust. Therefore, other factors in the engine gas path like the fuel flow, vibration, and the fan speed might affect [EGT](#). This project seeks to develop a robust machine learning predictive model for [EGT](#) prediction

by incorporating additional engine performance variables alongside the EGT and flight conditions such as altitude and flight speed. Such a robust predictive model could lead to enhanced and accurate [EGTM](#) predictions hence maximizing maintenance resources and the overall aircraft operational efficiency.

1.3 Research Aim

To develop a predictive maintenance model for an aircraft engine to reduce downtime, and optimize maintenance schedules.

1.4 Research Objectives

This research aimed to address the following objectives:

- (a) To identify the factors investigated in aircraft engine maintenance.
- (b) To develop a predictive maintenance model that forecasts Exhaust Gas Temperature margin.
- (c) To deploy the predictive model for maintenance monitoring and decision-making support.

1.5 Research Questions

The research questions addressed in this study were as follows:

- (a) Which engine parameters are critical in predicting EGT margin?
- (b) Which Machine learning algorithm effectively forecasts the EGT margin for predictive maintenance strategies?
- (c) How can the predictive maintenance model be effectively integrated into the maintenance workflow of aircraft engines?

1.6 Scope and Limitations of the Study

1.6.1 Scope

The study included the following key components:

- (a) Identification and selection of engine data.

- (b) Developing and deploying machine learning models for predictive maintenance using real-world engine data.
- (c) Model interpretability: SHAP analysis was used to explain the contribution of each input feature to the model's predictions.
- (d) The analysis is confined to the takeoff phase of flight, which is critical for understanding variations in Exhaust Gas Temperature (Exhaust Gas Temperature (EGT)) margin.

1.6.2 Limitations

Here are some of the limitations of this research work:

- (a) The findings are not representative of the entire CFM56 engine family, as they are based on one specific engine instance.
- (b) Model performance may be influenced by unknown or inconsistent maintenance practices during the data collection period.

1.7 Research Justification

In this project a data-driven predictive maintenance model that predicts EGT margin based on historical performance was developed. This model can allow for timely intervention, reduced downtime, and optimized maintenance schedules, enhancing the reliability of aircraft operations while reducing costs associated with unplanned maintenance actions. These outcomes align with the United Nations Sustainable Development Goals (UNSDG), particularly SDG 9: Industry, Innovation and Infrastructure. According to SDG 9, innovation and comprehensive infrastructure fuels inclusive and sustainable industrialization which could result in effective economic planning that facilitate efficient use of resources [United Nations \(2023\)](#).

The predictive model also contributes to the ICAO Global Aviation Safety Plan (GASP), which seeks to enhance aviation safety through continuous improvement [International Civil Aviation Organization \(ICAO\) \(2023\)](#). This is because the model is expected to enhance aviation operational safety by identifying and addressing potential issues before they escalate.

Chapter 2: Literature Review

2.1 Overview of the Operation of Aircraft Engines

Four major forces act on an aircraft in flight: drag, weight, lift, and thrust. To achieve flight, an aircraft must overcome its weight, which is done by generating an upward force, called lift, produced by the wings. As an airplane moves through the air, drag is generated due to air resistance to the plane's motion. Drag acts in the opposite direction as the flight; it resists the plane's motion. To counteract drag, aircraft use a propulsion system to generate thrust. Aircraft jet engines are crucial for flight operation because they generate the thrust needed to overcome drag to propel the plane forward [Hu et al. \(2010\)](#). Figure 2.1 shows the major forces of flight.

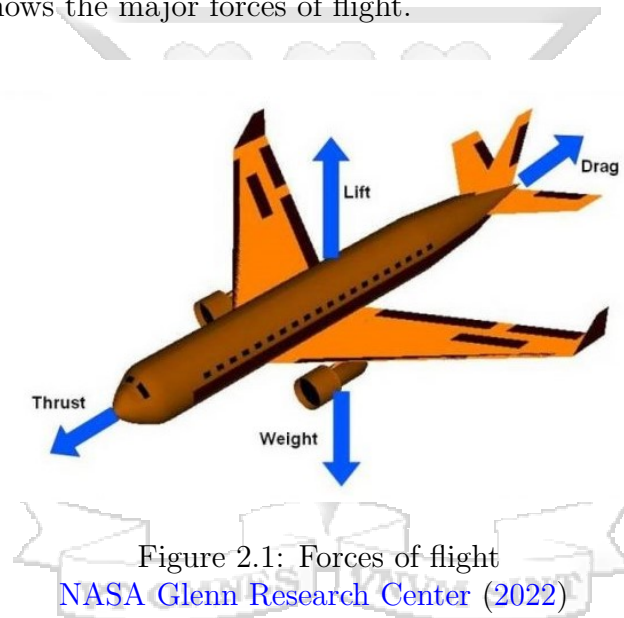


Figure 2.1: Forces of flight
[NASA Glenn Research Center \(2022\)](#)

The jet engine operates on the Brayton cycle, where air is sucked in via the intake. The air is then compressed in the compressor, increasing its pressure and temperature. It is mixed with jet fuel droplets and ignited, producing hot exhaust gases and a high increase in temperature. The high-temperature exhaust gases rush past turbine blades that extract energy from the gas, lowering the pressure and temperature. The turbine directs gas out of the engine through the exit nozzle at high velocity, generating thrust that propels the aircraft forward [Ehrich and Baxter \(2024\)](#). Figure 2.2 shows the major parts of a jet engine.

The temperature of the exhaust gases (EGT) is a vital indicator of the engine's overall health and performance. One major sign of engine failure is the rise in EGT, which leads




Figure 2.2: Major components of a Jet Engine
 Ahmed et al. (2015)

to turbine wear and a decline in remaining useful life. Hence, the EGT is a key parameter in forecasting, performance optimization, and monitoring of engine health Yildirim and Kurt (2018), Fedele et al. (2020). The maximum allowable EGT during take-off is called the EGT redline, while the EGT margin is the difference between the EGT redline and the actual EGT. Yildirim and Kurt (2016). EGT margin gives information on engine health conditions. Monitoring the EGT margin is critical in predictive maintenance as it signals the need for maintenance actions. Yildirim and Kurt (2018). Figure 2.3 illustrates the EGT margin.




Figure 2.3: Illustration of EGT margin
 Ahmed et al. (2015)

Aircrafts operate at high temperatures, pressures, and speeds; hence, they must be operated within specified limits to avoid structural failure Yildirim and Kurt (2018). During take-off, the EGT should be kept low to avoid exceeding the pre-defined limits Ackert (2011). That is why aircraft engines undergo scheduled and unscheduled maintenance checks at various intervals, depending on the stipulated calendar time, flight hours, or flight cycles, to ensure airworthiness and flight safety Mofokeng et al. (2020). Maintenance is critical due to the fault-tolerant nature of aircraft engines, where minor faults can lead to catastrophic failures if left unchecked. Hence, high standards of reliability

and efficiency are set for flight operations and aircraft maintenance programs, failure to which can lead to safety concerns, operational disruptions, and high maintenance costs [Ackert \(2010\)](#).

2.2 Related Works

The evolution of aircraft maintenance practices has shifted from Corrective to preventive, and eventually to predictive maintenance. Corrective maintenance is performed after a system failure occurs. Preventive maintenance is scheduled maintenance where a system is periodically repaired or replaced after a specific number of flight hours or calendar times. Predictive maintenance is where a system is continuously monitored, and maintenance is performed based on real-time and historical data, statistical inference, integrity factors, and machine learning approaches [Adryan and Sastra \(2021\)](#).

[EGT](#) prediction techniques based on real flight data are categorized into model-based and data-driven methods. Model-based methods rely on the mathematical and physical design of the system. Such methods require accurate simulation of the operation of mechanical systems, which might be challenging to achieve and require professional competence. Data-driven methods are more accessible to implement and hence commonly used in practice [Liu et al. \(2024\)](#).

Data-driven prediction techniques include statistical, machine, and deep learning techniques. Common statistical methods include the auto-regressive model, and the auto-regressive integrated moving average model. These methods are suitable for stationary and linear data; hence, they are not appropriate for [EGT](#) prediction, which changes often [Emer and Özbek \(2020\)](#). Increased sensor technology in aircraft has increased the application machine and deep learning approaches for predictive maintenance [Liu et al. \(2024\)](#). Several machine learning and deep learning techniques have been proposed for [EGT](#) prediction, as summarised below;

Yildirim and kurt used artificial neural network (ANN) in [EGT](#) prediction using real flight data. First, multiple regression was used to find the relationship between the target and predictor variables. In this case, the input variables were angle of attack, NI (fan) speed, pitch, N2 speed, roll, total air temperature, vertical acceleration, and ground/air variables, while the output was the [EGT](#). A t-test also revealed that all the predictor

variables were significant. Next, the artificial neural network technique was fitted on the predictor variables, where 11 different functions were used to predict [EGT](#). The results showed that the Levenberg-Marquardt backpropagation function was the best since it had the lowest MSE [Yildirim and Kurt \(2018\)](#).

In their research, Wang and Zhao fit four machine learning methods on real-world engine data to predict the exhaust gas temperature. The predictor variables were total atmospheric temperature, flight altitude, low and high pressure shaft speeds, fuel flow, and Mach number. The data was first filtered using an Unscented Kalman Filter to reduce noise. After performing a correlation analysis, the data was evaluated using four machine learning models; random forest, support vector regression (SVR), Generalized Regression Neural Network (GRNN), and radial basis neural network (RBF). The results showed that SVR and GRNN had higher prediction accuracy. The GRNN had the least RMSE and MAE, and used the least CPU time and the highest calculation efficiency [Wang and Zhao \(2023\)](#).

The Generalized Additive Model (GAM) was used on simulated turbofan engine data developed by National Aeronautics and Space Administration ([NASA](#)) to predict [EGT](#). The data had three subsets to explore the model's accuracy at different levels of deterioration. Subset one had one single component (HPT efficiency) deterioration; subset two had three components (High Pressure Turbine ([HPT](#)) and Low Pressure Turbine ([LPT](#)) efficiency, and mass flow capacity), while subset three had deterioration in all gas path components in terms of both efficiency and mass flow capacity. In all the subsets, the input parameters were physical core speed, fuel flow, static pressure at the High Pressure Compressor ([HPC](#)) outlet, total temperature at the [HPC](#) outlet, total pressure at the Low Pressure Compressor ([LPC](#)) outlet, and total temperature at the [LPC](#) outlet. These features were selected before they strongly correlated with the target variable based on the Pearson correlation matrix. RMSE and R squared were used for model evaluation with good desirable results. The GAM showed good interpretability and high prediction accuracy in all deterioration levels [Apostolidis et al. \(2022\)](#).

Ullah et al fit a three-layer LSTM model with an Adam optimizer was used on actual flight data for [EGT](#) prediction. The model used turbo compressor rotor rotation frequency, fuel feeding, inlet air temperature, fan rotation frequency, inlet air Pressure, gas temperature

after the fan turbine, nozzle throat, air pressure after the compressor, gas pressure after the fan turbine, throttle position, oil temperature going out the engine, pressure difference on the Fuel Filter, and fuel temperature after the fuel filter as input variables and **EGT** as the target variable. Two sets of in-flight data were used: data from engines that showed no faults during regular inspections (healthy data) and data from engine that exhibits faults such as nicks, faults and cracks (faulty data). Healthy data was used for model training, and faulty data was used for testing and validation. A mean absolute error of 5.7°C was achieved [Ullah et al. \(2023\)](#).

In their study, Kefalas et al. used symbolic regression (SR) on real-world engine data to discover the relationship between **EGT** and other engine parameters. SR was used because some machine learning models rely heavily on assumptions such as stationery, and other models, such as neural networks, are "black boxes." As such, their work is viewed in relation to inputs and outputs, making it challenging to interpret. The data used related to a GenX engine recorded during the cruising phase of four flights. The input variables were **HPT** clearance, gas temperature at the inlet of **HPC**, corrected fan speed, load control of engine generators, SV Minimum Main Fuel Split Regulator, Selected Variable Bleed Valve Position, and non-dimensional fuel flow. The experimental results showed good model accuracy and were easy to interpret with a MAE of 3°C [Kefalas et al. \(2022\)](#).

An Enhanced Scale-Aware efficient Transformer (ESAE-Transformer) was used in **EGT** prediction. The model optimized computational efficiency and identified key temporal patterns by merging the standard transformer architecture with the Multi-Head Prob-Sparse Attention mechanism. A Multi-Scale Feature Aggregation Module was included to improve the handling of complex temporal features, thereby boosting the model's predictive performance. Real aircraft flight data recorded in climb, cruise, and landing phases was used. Predictor variables used were Mach number, flight altitude, power lever angle, fuel flow rate, ambient temperature, physical fan speed, and physical core speed. The model used the AdaX optimization technique, and achieved a mean absolute **EGT** prediction error of 3.47°R , which aligns with the industry standards [Liu et al. \(2024\)](#).

A novel hybrid **EGT** prediction method, which combined the Nonlinear AutoRegressive with exogenous inputs (**NARX**) and Moving Average, was proposed by Ma et al [Ma](#)

et al. (2023). The FAE-LSTM model was used to build the NARX for a rough EGT estimation, after which a vanilla Long-Short-Term Memory (LSTM) is used to build the Moving Average model to determine the prediction error. Real aircraft flight data recorded in climb, cruise, and landing phases was used. Input variables included flight altitude, ambient temperature, Mach number, fuel flow, power lever angle, fan speed, and core speed. Compared to the individual baseline models, the hybrid model reduced the Mean Absolute Error (MAE) and Root Mean Square Error (RMSE) by at least 18.47% and 13.23%, respectively.

2.3 Research Gap

Most studies focus on predicting Exhaust Gas Temperature (EGT) directly, but EGT alone does not provide sufficient insight into engine health. A more critical metric is EGT Margin — EGT Hot Day Margin in high Outside Air Temperature (OAT) conditions—because it quantifies how much buffer remains before the engine reaches its operational limit. EGT Margin is defined as the difference between the maximum allowable EGT and the actual EGT at takeoff, directly indicating an engine’s structural integrity. A declining EGT Margin suggests thermal stress, wear, and potential material degradation, which can compromise engine performance and safety over time.

Additionally, some studies such as Ullah et al. relied solely on engine operational history, thereby overlooking other potentially valuable data inputs such as environmental conditions that could enhance model accuracy and reliability Ullah et al. (2023). Furthermore, some studies such as Apostolidis et al., utilized simulated data rather than real-world datasets. While simulations can offer controlled environments for experimentation, they may not fully capture the complexity and variability of actual engine behavior, thus limiting the applicability of their findings in real-world scenarios Apostolidis et al. (2022). To address the concern of machine learning models being perceived as ‘black boxes’, as noted by Kefalas et al. Kefalas et al. (2022), this study incorporates SHapley Additive Explanations (SHAP) analysis to enhance model interpretability. By doing so, the contribution of each input feature to the model’s predictions can be visualized and better understood, thereby increasing transparency and trust in the model’s decision-making process.

Aircraft maintenance teams can rely on [EGT](#) Margin trends for engine condition monitoring and timely maintenance scheduling. While various machine learning models have been applied to [EGT](#) prediction, limited research has focused on [EGT](#) Margin prediction, which depends on multiple operational and environmental factors. Predicting the [EGT](#) Margin requires incorporating variables such as take-off airport altitude, ambient temperature, engine age, and operational history. Accurate [EGT](#) Margin predictions can help airlines assess remaining engine life, optimize maintenance schedules, and reduce operational costs. This study aims to fill this gap by developing a machine learning model specifically for predicting [EGT](#) Margins, enhancing proactive maintenance, structural health monitoring, and safety in aircraft maintenance operations.



Chapter 3: Methodology

This study adopted the Cross-Industry Standard Process for Data Mining (CRISP-DM), a widely recognized and iterative framework commonly used in data mining projects, as illustrated in Figure 3.1.




Figure 3.1: CRISP-DM framework

1. **Business Understanding:** This initial phase involves defining the core problem, which may stem from business goals or research questions. It includes engaging with stakeholders to ensure that the project objectives are clearly aligned with their expectations.
2. **Data Understanding:** At this stage, relevant data is collected and assessed for its suitability to address the research goals. This helps establish a clear understanding of the dataset's characteristics and limitations.
3. **Data Preparation:** This step includes cleaning, transforming, and engineering fea-

tures from the raw data. Collaboration with subject matter experts is essential here to ensure the dataset is comprehensive and ready for modeling.

4. **Modeling:** During this phase, various modeling techniques are explored and applied to the data. Emphasis is placed on selecting a diverse set of algorithms, which are then trained to perform predictive tasks.
5. **Evaluation:** The models are rigorously assessed based on performance metrics such as accuracy and class discrimination, ensuring they meet the predefined research or business objectives.
6. **Deployment:** The most effective model is integrated into a practical setting, such as a web application or dashboard. This allows the model to serve its intended function, with ongoing monitoring to maintain its performance and relevance.

3.1 Business Understanding

The business understanding phase of this project focused on developing a system capable of accurately predicting [EGTM](#) in aircraft engines, with the goal of enhancing maintenance planning and operational efficiency in the aviation sector. By addressing the need for reliable EGT margin forecasts, the project aims to contribute to early detection of engine performance anomalies and potentially reduce maintenance costs and downtime. This section included the following subsections:

1. **Identifying Stakeholders:** The main beneficiaries of this system include airline maintenance engineers and aircraft fleet operators who rely on accurate performance metrics to make informed decisions. While this project does not directly engage these stakeholders, their needs and typical workflows informed the design of the EGT prediction system. A user-friendly web portal was developed to simulate how such stakeholders might interact with predictive insights in an operational setting.
2. **Defining Objectives:** The core objective of this project was to build a deep learning model capable of predicting [EGTM](#) using historical engine sensor data. Specific goals included evaluating the model's performance using metrics such as MAE, MSE, RMSE, R^2 , and adjusted R^2 , deploying the best-performing model on a web platform, and offering users a straightforward interface for uploading data

and receiving predictions.

3. **Determining Requirements:** Requirements for the system were driven by the need to support diverse data inputs (both manual entry and file upload), ensure rapid and accurate predictions, and enable users to receive notifications of results.

3.2 Data Understanding

Aircraft engines are equipped with sensors that continuously monitor various performance metrics such as fuel flow, temperature, oil levels, pressure, vibration, and performance metrics. This data is collected by the aircraft's onboard systems, such as the Flight Data Acquisition Unit. The data is transmitted via Aircraft Communications Addressing and Reporting System ([ACARS](#)), a digital data link system used in aviation to relay data between aircraft in flight and ground stations. [ACARS](#) also makes aircraft engine data visible to the Original Equipment Manufacturer ([OEM](#)) by facilitating real-time or near-real-time data transmission. The dataset to be used in this project is obtained directly from the OEM site.

The dataset, comprising 10,133 records and 24 variables, was collected between January 2016 and September 2024, and specifically pertains to the take-off flight phase. The dataset comprise columns that contain general information on the specific aircraft and engine, such as airline name, aircraft serial number, flight date-time, engine position, aircraft registration number, engine family, engine type, engine installation date, flight phase, and aircraft type. The data also include numeric columns that contain engine performance variables as explained in [Table 3.1](#) below.

Each of these engine parameters influence engine performance and can cause engine failure in the following ways:

1. **EGT and EGT Margin**

Elevated [EGT](#) often indicate increased thermal stress on engine components, particularly the turbine blades, which can lead to accelerated wear or failure if temperatures exceed design limits. The [EGT](#) margin represents the difference between the maximum permissible [EGT](#) and the observed [EGT](#). A decreasing [EGT](#) margin is indicative of engine deterioration and impending need for maintenance [Kurz and Brun \(2001\)](#).

Table 3.1: A description of engine performance parameters
fed (2012)

No	Parameter Name	Description	Units
1	EGT Margin	This is the temperature difference between the incurred takeoff EGT and the Redline (maximum) EGT.	°C
2	EGT	This is the temperature of the exhaust gases.	°C
3	Mach	The ratio of the aircraft's speed to the speed of sound.	-
4	Fuel Flow	This is the rate of flow of jet fuel to the engine for combustion to take place. A desirable fuel flow rate ensures sufficient supply of fuel to the engine for combustion.	pph (pounds per hour)
5	Vibration N1#1 Bearing	A measure of the oscillations at the N1 stage of the engine at bearing 1.	Aircraft units
6	Vibration N2#1 Bearing	A measure of the oscillations at the N2 stage of the engine at bearing 1.	Aircraft units
7	Oil Temperature	This is the temperature of engine oil as it flows through the engine to cool the mechanical components of the engine.	°C
8	Oil Temperature Smoothed	This is the moderated temperature of engine oil as it flows through the engine to cool the mechanical components of the engine.	°C
9	Oil Pressure	This is a measure of the pressure of engine oil as it flows through the engine to cool mechanical components of the engine.	PSIG (Pounds per Square Inch Gauge)
10	Total Air Temperature	A measure of the static air temperature and the rise in air temperature as the aircraft moves through it.	°C
11	Indicated Fan Speed	The rotational speed of the fan.	RPM(revolutions per minute)
12	Thrust Derate	The percentage reduction from the maximum rated thrust	
13	Thrust Derate Smoothed	The percentage gradual reduction from the maximum rated thrust	
14	Core Speed	The rotational speed of the high-pressure spool section of the engine, measured at the high-pressure compressor.	RPM (revolutions per minute)
15	Altitude	The height that an aircraft flies above sea level.	Feet (ft)

2. Mach

Mach affects engine inlet conditions, especially in supersonic flight, altering the air-flow dynamics into the compressor. Incorrect Mach number management can result in shock waves at the intake or compressor stalls, affecting performance [Mattingly \(2002\)](#).

3. Fuel Flow

Fuel flow determines the amount of fuel supplied to the engine. Higher-than-normal fuel flow rates can be symptomatic of degraded engine efficiency or internal leaks. Conversely, restricted fuel flow can cause power loss or flameout, which may compromise flight safety [Daggett et al. \(2003\)](#).

4. Vibration (N11 and N21 Bearings)

Vibrations in the N1 and N2 rotors (low and high-pressure spools, respectively) can indicate imbalances or mechanical wear in bearings. Excessive vibrations can lead to fatigue failure in rotating components and are a leading indicator of bearing or blade damage [Fabry and Češkovič \(2017\)](#).

5. Oil Temperature

Oil temperature is critical for maintaining engine lubrication quality. Elevated oil temperatures may degrade oil viscosity, reducing its ability to lubricate and protect engine components from friction-induced damage. Over time, this can lead to bearing failure and rotor seizure [Boyce \(2011\)](#).

6. Oil Pressure

Oil pressure must be maintained within specific limits to ensure consistent lubrication across engine systems. Low oil pressure may suggest leaks, pump failure, or clogged filters, leading to inadequate lubrication and potential engine seizure [Kurz and Brun \(2001\)](#).

7. Total Air Temperature (TAT)

TAT combines ambient and kinetic heating of air entering the engine. It influences air density and, consequently, compressor performance. High TAT can reduce thrust output and contribute to compressor overheating [Mattingly \(2002\)](#).

8. Indicated Fan Speed

Indicate fan speed is directly related to engine thrust generation. Deviations from expected the specified values may indicate performance deterioration, fan blade damage, or operational anomalies [Daggett et al. \(2003\)](#).

9. Thrust Derate

Thrust derating is the intentional reduction of engine thrust to prolong engine life and reduce maintenance costs. Improper derating or excessive thrust usage can accelerate wear and lead to thermal degradation of engine components [Boyce \(2011\)](#).

10. Core Speed (N2)

N2 is the rotational speed of the high-pressure spool. Like N1, it provides insights into core engine performance. Abnormal N2 readings can point to compressor or turbine inefficiencies, potentially leading to core degradation or failure [Matthaiou et al. \(2017\)](#).

11. Altitude

Altitude affects ambient air pressure and temperature, both of which impact engine efficiency and thrust. At higher altitudes, thinner air reduces engine intake mass flow, possibly leading to compressor stalls or insufficient combustion under certain conditions [Mattingly \(2002\)](#).

3.3 Data Preparation

3.3.1 Handling Missing Values

Null entries, in this case, could result from sensor malfunction, no record being made, or from consolidating data from different time frames. All numerical columns describing engine performance have missing values. For each case, the missing values are less than 5% of the entries hence were dropped. This is in accordance with an article by Krishna Sree who recommends that for rows, If less than 5% of rows have missing values in a particular column, then the missing values can be dropped in those rows [Sree \(2024\)](#).

3.3.2 Feature Engineering

To take in to account engine wear and tear, a new column was created to show how long the engine had been in operation. The column, named "Days Since Install" was created by getting the difference between the "engine install date" and the "flight datetime" columns. The value obtained was then converted to days.

3.3.3 Dealing with Outliers

Outliers are data points that are way above or below the average range of values. Since numerical fields in the dataset had different ranges of values, scaling using the Min-max scaling from the Scikit-learn machine-learning library had to be done to bring all the numerical columns to the same value scale [Pedregosa et al. \(2011\)](#). The Min-Max scaling technique was used because it is appropriate when features have different ranges or units which is true for the dataset used in this project [Sahin \(2024\)](#). The Min-Max scaling formula is used to transform features to lie within a specific range, typically between 0 and 1 using this formula

$$X_{\text{scaled}} = \frac{X - X_{\text{min}}}{X_{\text{max}} - X_{\text{min}}} \quad (1)$$

Where:

- X is the original value to be scaled.
- X_{min} is the minimum value of the feature.
- X_{max} is the maximum value of the feature.
- X_{scaled} is the scaled value.

3.4 Machine Learning Modeling

Five regression models — Random Forest, Support Vector Regressor (SVR), K-Nearest Neighbors (KNN), Decision Tree and Linear Regression — are used to predict the EGT Hot Day Margin of aircraft engines. Each model follows a supervised learning architecture where the dataset is split into training (80%) and testing (20%) subsets. These models are inbuilt and available in the Scikit-learn Python library [Pedregosa et al. \(2011\)](#).

Random Forest is an ensemble method that merges the outputs of multiple decision trees to a single output. Each tree is constructed using a subset of the training data and features, introducing diversity among the trees and reducing the risk of overfitting. For this project, the Random Forest model was set up with 50 estimators and a maximum depth of 3.

Support vector regressor is a supervised machine learning algorithm that fits a hyperplane, that minimizes the difference between the actual and predicted values. In this project, Support Vector Regressor was set up with all its default hyperparameters including RBF kernel, an epsilon value of 0.1, and regularization parameter set to 1.0.

K-Nearest Neighbors is a non-parametric model that predicts the target value for a given input by averaging the values of its 'k' nearest neighbors in the feature space. The separation distance between data points is usually calculated using Euclidean distance, although other distance metrics can be used. In this project, KNN regressor was configured with default hyperparameters, that is, five neighbors.

Decision Tree regressor is a non-linear model that uses a tree-like model to make predictions. It splits the data into smaller subsets based on decision rules derived from the input features. Each split is made to minimize the error in predicting the target variable. Decision Tree regressor was set up with a maximum depth of 3.

The linear regression model captures the linear relationship between the target and predictor variables by fitting a linear equation from the training data, resulting in line of best-fit. The slope of the line indicates how much the predictor variable changes for a unit change in the target variable. Linear regression model was set up with its default parameters.

For the deep learning model, a fully connected feed-forward neural network was build to predict EGT Hot Day margin. The architecture consisted of an input layer, four hidden layers, and a final output layer. The hidden layers gradually decrease in size with the first containing 64 neurons, followed by subsequent layers with 32, 16, and 8 neurons, respectively. Each hidden layer utilized the ReLU activation function to introduce non-linearity. The model was compiled Adam optimizer, and the mean squared error is used as the loss function. Training is conducted over 100 epochs with a batch size of 32,

meaning the model updates its weights after processing 32 samples. A 20% validation split is used to monitor the model's performance on unseen data during training.

The deep learning model turned out to be the best-performing model. A grid search validation was set up to find the best optimizer and learning rate. Using these best hyperparameters, the model was fine-tuned by incorporating a 20% Dropout rate in each hidden layer and applying Early Stopping to prevent overfitting.

3.5 Performance Evaluation Criteria

3.5.1 Regression Models Performance Evaluation

Evaluation metrics were selected to evaluate the performance of the regression algorithms. The metrics were, mean absolute error, mean square error, root mean squared error, R-squared, and adjusted R-squared, described in the equations 2- 6 below. In the equations, n is the number of observations, y_i represents the actual value for the i -th observation, and \hat{y}_i represents the predicted value for the i -th observation [Hackeling \(2017\)](#), [Chicco et al. \(2021\)](#).

The best algorithm is the one that performs the best across the most relevant metrics. Typically, lower MSE, RMSE, and MAE values indicate better performance, and higher R^2 and Adjusted R^2 Scores indicate better model fit [Tatachar \(2021\)](#).

The Mean Absolute Error is defined as the mean of the absolute differences between the predicted values and the actual values over all the observations.

$$\text{MAE} = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i| \quad (2)$$

The Mean Squared Error measures the average of the squares of the errors, that is, the average squared difference between actual and predicted values.

$$\text{MSE} = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2 \quad (3)$$

The RMSE is the square root of the mean of the squared differences between actual and predicted values.

$$\text{RMSE} = \sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2} \quad (4)$$

The R^2 (Coefficient of Determination) score measures the proportion of variance in the dependent variable that is predictable from the independent variable(s).

$$R^2 = 1 - \frac{\sum_{i=1}^n (y_i - \hat{y}_i)^2}{\sum_{i=1}^n (y_i - \bar{y})^2} \quad (5)$$

The adjusted R^2 accounts for the number of predictors in the model and adjusts the R^2 score accordingly.

$$R_{\text{adj}}^2 = 1 - \left(\frac{1 - R^2}{n - k - 1} \right) \times (n - 1) \quad (6)$$

3.5.2 Model vs. Operational Limits Comparison

Model validation was conducted using the EGT Hot Day Margin operational limits specified in the CFM56-7B26 engine operational limits guide. According to the guide, the initial EGT Margin limit for a new engine is 60°C, while for a restored engine, the margin ranges from a minimum of 42°C to a maximum of 48°C [CFM \(2008\)](#). For this project, the internal EGT Hot Day Margins of the airline that provided the data were used, with values ranging from a minimum of 10°C to a maximum of 55°C. Any values falling outside this range were classified as 1 (indicating regions requiring a maintenance check), while values within the specified limits were labeled as 0 (showing within the operational range).

Model evaluation for comparing the predicted EGT Margin against operational limits was performed using classification metrics, including recall, F1-score, precision, and accuracy, as defined in Equations 6 to 10 below [Géron \(2022a\)](#). In the equations:

- (i) TP (True Positives): The number of correctly predicted instances where the EGT Margin remains within the acceptable operational limits.

- (ii) TN (True Negatives): The number of correctly predicted instances where the EGT Margin falls outside the operational limits.
- (iii) FP (False Positives): The number of instances where the model incorrectly predicted the EGT Margin to be within the acceptable range when it was actually outside the limits.
- (iv) FN (False Negatives): The number of instances where the model incorrectly predicted the EGT Margin to be outside the acceptable range when it was actually within limits.

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \quad (6)$$

$$\text{Precision} = \frac{TP}{TP + FP} \quad (7)$$

$$\text{Recall} = \frac{TP}{TP + FN} \quad (8)$$

$$\text{F1-Score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \quad (9)$$

$$\text{F1-Score} = \frac{2 \times TP}{2 \times TP + FP + FN} \quad (10)$$

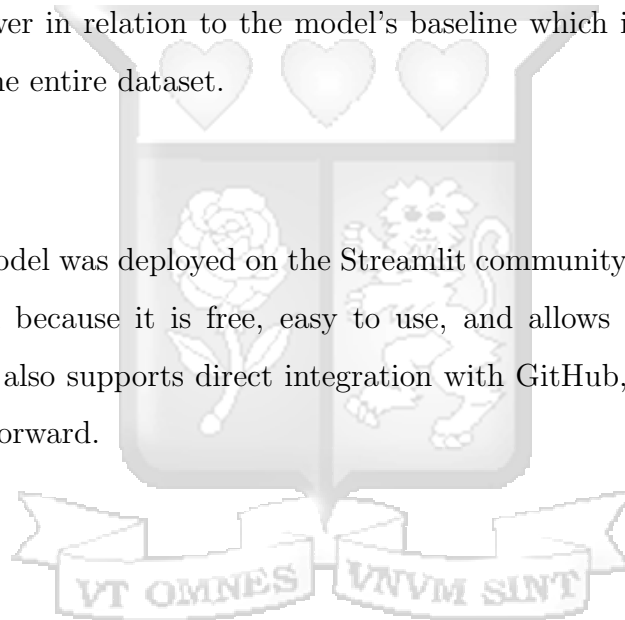
3.5.3 Shapley Additive Explanations (SHAP) Analysis

SHAP is a technique used to interpret machine learning models by demonstrating the impact of each feature on the final prediction, the relative importance of each feature compared to others, and the extent to which the model depends on feature interactions. In summary, it provides to provide explanations for specific predictions [Lundberg and Lee \(2017\)](#). [SHAP](#) analysis was done for all the models, and [SHAP](#) values were used to generate visualizations that depict how each feature impacts the final prediction through the following plots:

- (i) Summary Plot: Displays the distribution of **SHAP** values for all features, highlighting their overall impact on predictions.
- (ii) Summary Bar Plot (Feature Importance Plot): Ranks features based on their average absolute **SHAP** values, showing their relative importance.
- (iii) Waterfall Plot: Breaks down an individual prediction by visualizing how each feature contributes positively or negatively to the final prediction.
- (iv) Decision Plot: Shows how **SHAP** values accumulate across different features for a given prediction, showing the decision path.
- (v) Force plot: shows how each feature contributes to the final prediction, either pushing it higher or lower in relation to the model's baseline which is usually the average prediction of the entire dataset.

3.6 Deployment

The deep learning model was deployed on the Streamlit community cloud platform. This platform was chosen because it is free, easy to use, and allows quick sharing of web apps with others. It also supports direct integration with GitHub, making updates and deployment straightforward.



Chapter 4: System Design and Architecture

A web-based predictive maintenance application was developed for monitoring aircraft engine health through Exhaust gas temperature margin analysis using the deep learning model as the backend. It provides automated alerts when predicted values exceed operational thresholds, enhancing early detection of potential failures.

The system was implemented using Visual Studio Code, version-controlled via GitHub, and deployed on the Streamlit Community Cloud. This chapter details the design architecture of the app.

4.1 System Modeling

This study utilized the Unified Modeling Language Unified Modeling Languages (UML) for system modeling, a widely recognized standard in software engineering known for its ability to effectively represent and document system architecture. UML provides a common visual language that facilitates clear communication and understanding among system architects, developers, and other stakeholders.(Koc et al., 2021).

Figure 4.1 illustrates the relationship between various components in the system.

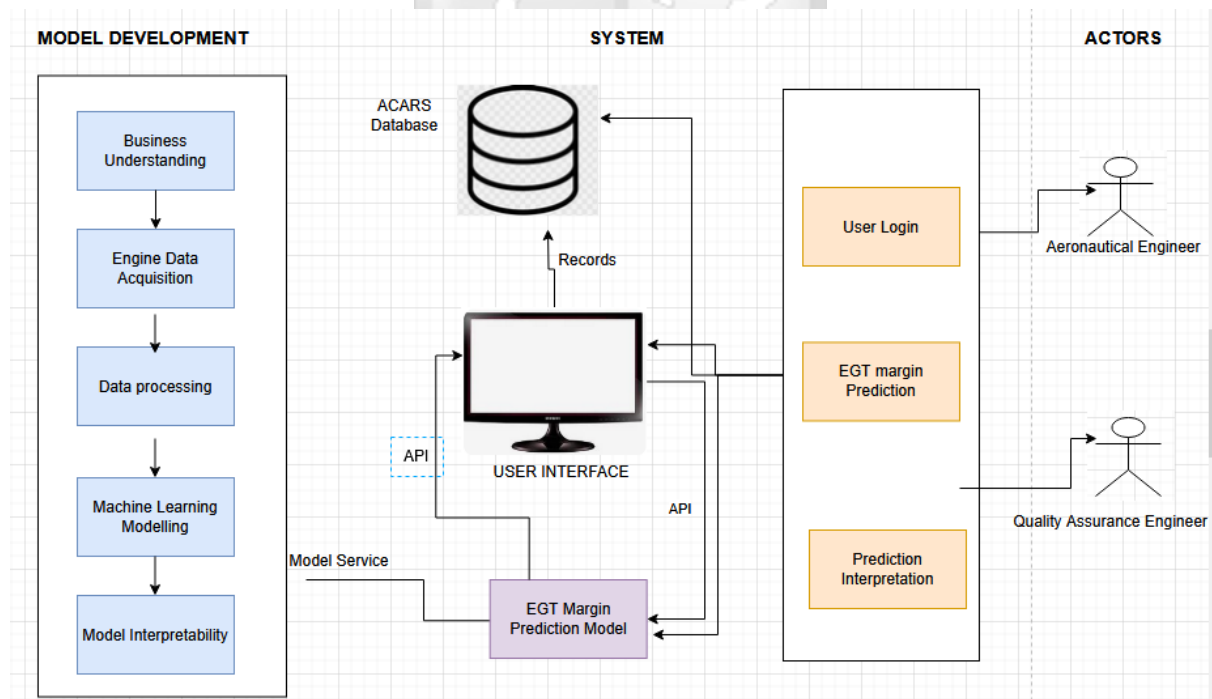


Figure 4.1: UML diagram

4.2 System Components

In this study, the system comprised of two major components: The web portal and database. The database serves as a temporary data repository, storing data upload on the app. This data includes input data uploaded the user, email alert data, and user login.

4.2.1 Database

The [ERD](#) illustrates the temporary data flow and relationships within your Streamlit-based EGT margin prediction application. Although no actual database is created, this diagram represents how data is structured and temporarily handled in-memory during a user session. The structure comprise of the following components:

1. User session: Represents a user's interaction session with the app, which begins when they access the application.
2. Input data: Captures the engine parameters entered manually or uploaded via file. Each session can have multiple input records.
3. Prediction result: Stores the model's prediction and related performance metrics for each input data record.
4. Notification: Records details if the user chooses to send an email notification after a prediction is made.

These components interact in a one-to-many or one-to-one fashion to reflect a typical usage workflow as shown in [Figure 4.2](#).

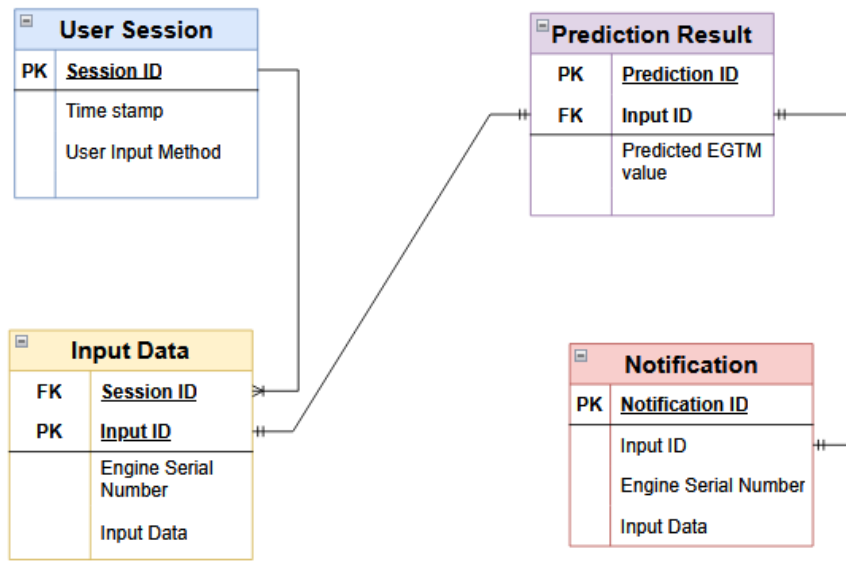


Figure 4.2: ERD of the tables

Table 4.1 details the structure of temporary data storage entities used in the EGT prediction web application.

Table 4.1: Database Tables

Entity	Attribute	Data Type	Description
User Session	session id timestamp user input method	int Datetime char(10)	Unique ID for the user's session When the session started Indicates whether data was uploaded or entered manually
Input Data	input id session id engine serial number input data	int int char(20) int, csv, xlxs	Unique ID for the input data Links the input to a session Unique identifier of the engine Data that a user uploaded or entered manually
Prediction Result	prediction id input id predicted EGT margin	int int float	Unique ID for the prediction result Links prediction to input data The predicted EGT margin value
Notification	notification id prediction id email address timestamp message sent	int int char(50) Datetime char(3)	Unique ID for the email notification Links notification to a prediction Recipient of the notification When the notification was sent Flag indicating if message was successfully sent

4.2.2 Web portal

The web portal was structured into three main pages: the Overview Page, the Predictions Page, and the Interpretation Page. The Overview Page served as an introductory guide for users, providing a clear explanation of the application's functionality. It detailed the workflow from data input to prediction output and included guidance on how to interpret various plots and visualizations presented throughout the portal.

The Predictions Page enabled users to interact with the predictive maintenance model by either uploading engine data files or manually entering relevant parameters. Upon submission, the system processed the input and returned Exhaust Gas Temperature (EGT) margin predictions along with other key indicators.

The Interpretation Page focused on model transparency and explainability. It featured [SHAP](#) waterfall plot, which illustrated the contribution of each input feature to individual predictions, helping users to understand the factors influencing model decisions.

1. Site Map

A site map is a blueprint of a website that illustrates how its various pages are organized and connected. It serves as a navigational and planning tool, both for users and for those involved in the design, development, and maintenance of the website as shown in [Figure 4.3](#)

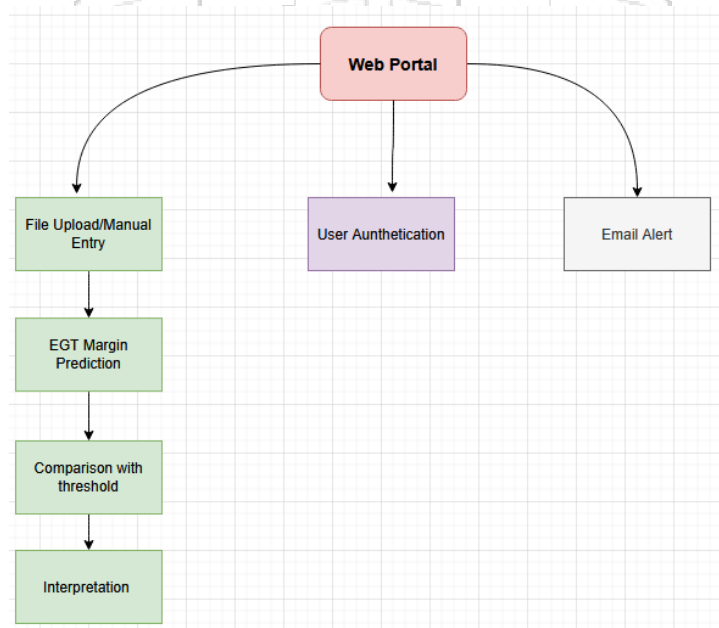


Figure 4.3: Sitemap of the web portal

2. Wireframes

Wireframes are basic, visual outlines of a website or application's layout that show where key components like navigation bars, content areas, and interactive elements will be positioned. They provide a clear framework for how the interface will function. By offering an early visual guide, wireframes help teams plan and refine the user experience before moving into detailed design or development. The drawio app (<https://www.drawio.com>) was used to create the wireframes of the web app pages.

a. Overview page wireframe

Figure 4.4 shows the wireframe of the overview page

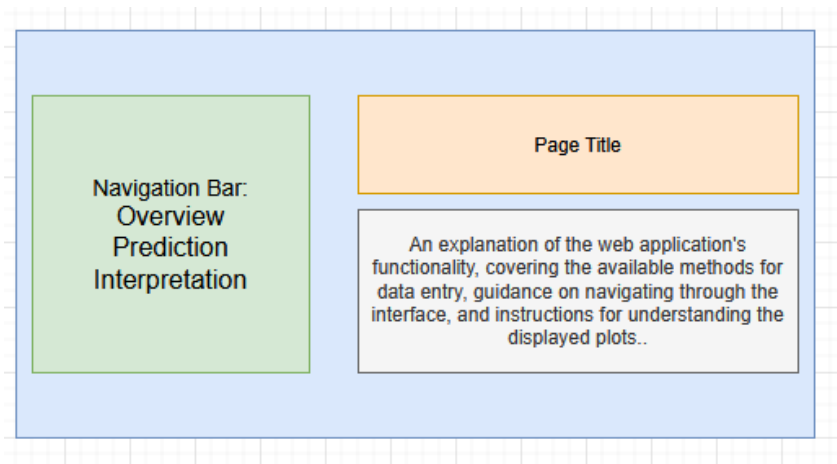


Figure 4.4: Overview page wireframe

b. Predictions wireframe

User have an option of either upload a file or manually entering the data as depicted in Figure 4.5

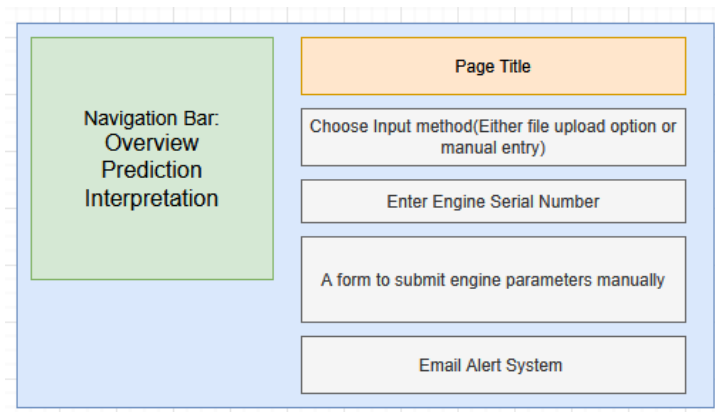


Figure 4.5: Predictions page wireframe

c. Interpretation wireframe

In this section, the app displays a [SHAP](#) waterfall plot that explains how each feature contributed to the final decision as shown in [Figure 4.6](#)

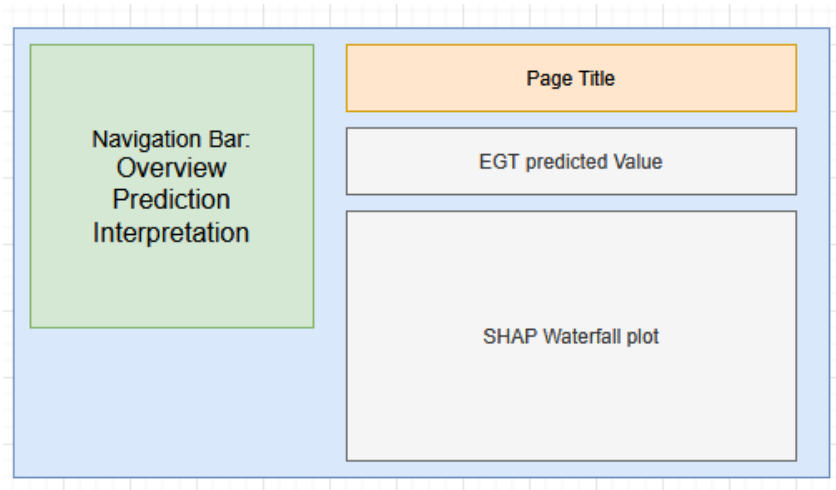



Figure 4.6: Interpretation page wireframe



Chapter 5: System Implementation and Testing

This chapter details the development and deployment of EGT margin prediction system, highlighting its key components and architecture. It also outlines the rigorous testing procedures undertaken to validate the system's functionality, usability, compatibility, and security

5.1 System Implementation

5.1.1 Database

In this research, the primary source of engine performance data is [ACARS](#), a digital data-link system widely used in aviation to transmit short messages between aircraft and ground stations. These messages contain key operational and maintenance data collected from onboard sensors and systems during various flight phases. In real-world operations, accessing [ACARS](#) data typically involves an authentication process, after which the user selects the specific engine model, engine type identified by serial number, relevant parameters, and the desired time period for data retrieval.

However, for the purposes of this project, direct interaction with the [ACARS](#) system was not undertaken. Instead, a pre-existing dataset derived from [ACARS](#) messages was provided. As such, data uploaded or manually entered by users in the app is temporarily stored on the Streamlit community cloud platform to facilitate processing and prediction within the web portal as illustrated in the [ERD](#) in Figure 4.2. When a user logs out, this temporary data is erased to ensure privacy and data security.

5.1.2 Web portal

The web portal was implemented using Streamlit community cloud platform and it comprises the following core components: Firstly, the app has restricted access with a Login page, requiring a username and password before accessing the webpage. The app then consists of three pages:

1. Overview Page

This is where a user learns what the app does and how to navigate the app effectively. It provides a guided walkthrough of the app's features, ensuring users

understand how to upload data, interpret predictions, and set up alerts. Additionally, it highlights key functionalities that enhance user experience and streamline interaction with the app as shown in Figure 5.1.

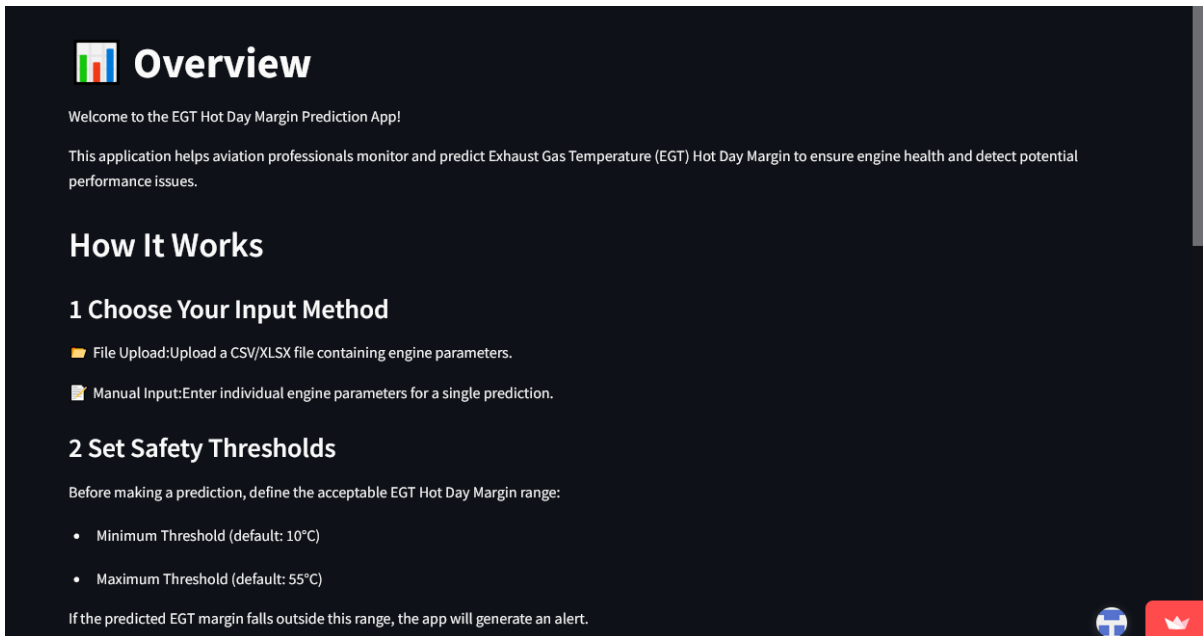


Figure 5.1: Overview page

2. Predictions page

This is where the user makes predictions based on the data they have. These predictions help in identifying trends, patterns, or anomalies within the dataset. The insights gained can be used to make informed decisions and optimize future outcomes. The deep learning model, which was the best-performing algorithm, generates the predictions. The app offers two input methods:

a) **File Upload Method:**

- (i) Users can upload a CSV or Excel file containing historical engine performance data.
- (ii) The app automatically processes the data, checking whether the EGT Hot Day Margin is already included or needs to be predicted.
- (iii) If the EGT Hot Day Margin is included in the uploaded data, the app acts as an automated system of checking whether the value exceeds thresholds and sends alerts to a set email address.

- (iv) If predictions are required, the app uses the deep learning model to generate EGT values.
- (v) The predicted values are compared against predefined safety thresholds to detect anomalies.
- (vi) The email alert highlights values that exceed safe limits, making it easy for users to spot potential issues.

b) **Manual Entry Mode:**

- (i) Users manually input individual engine parameters such as Mach, Fuel Flow, Vibration, and Oil Temperature.
- (ii) The app predicts the EGT Hot Day Margin based on the provided inputs.
- (iii) The system provides immediate feedback on whether the engine is operating within a safe range or requires attention.

Figure 5.2 illustrates the predictions page, presenting the uploaded data alongside the resulting predicted EGT margin value.

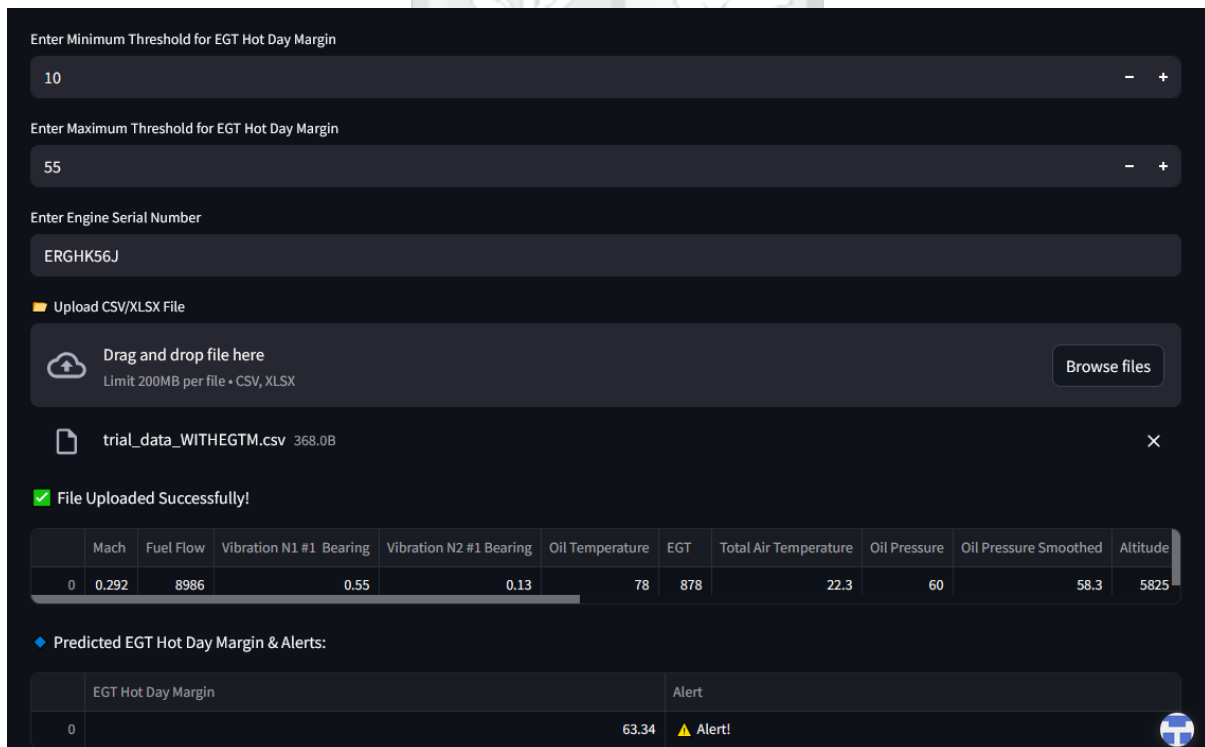


Figure 5.2: Predictions page

3. Interpretation Page

This page shows a SHAP waterfall plot that helps explain why a specific prediction was made for a specific input row of data. The visual starts with a base value and sequentially adds or subtracts SHAP values of each feature, showing their positive or negative impact on the final prediction as shown in Figure 5.3

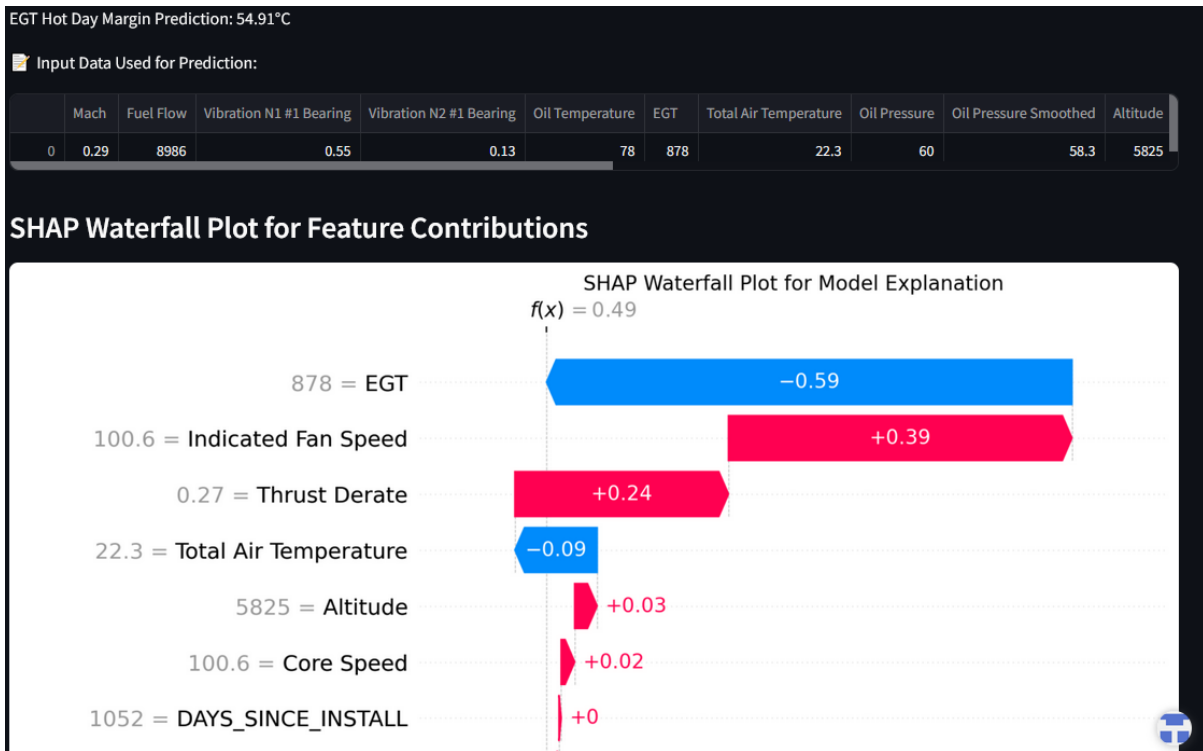


Figure 5.3: Interpretation page

Finally, there is a log-out option that resets the app back to its initial state, clearing all user inputs, uploaded files, and session data to ensure a fresh start for the next use.

The web app serves two major purposes:

1. Automation: The app allows users to upload take-off data, including EGT Margin, and automatically checks it against predefined thresholds. If an anomaly is detected, it can send an email alert to a designated recipient. This option allows for near-real time monitoring.
2. Flight Planning: For input data without EGT Hot Day Margin values, the app predicts the EGT Margins and assesses whether they exceed operational limits. If the predicted margins cross the threshold, preventive maintenance actions can be

planned in advance to mitigate potential engine performance issues.

5.1.3 Email Notification

After generating predictions on the predictions page, users have the option to send an email notification to a specified address. This feature enables them to conveniently share the results of the model's output, including key values and insights, with relevant stakeholders or team members. It enhances communication and supports timely decision-making based on the predicted data. The notification conveys details including the engine serial number, the defined threshold, and the predicted value shown in Figure 5.4

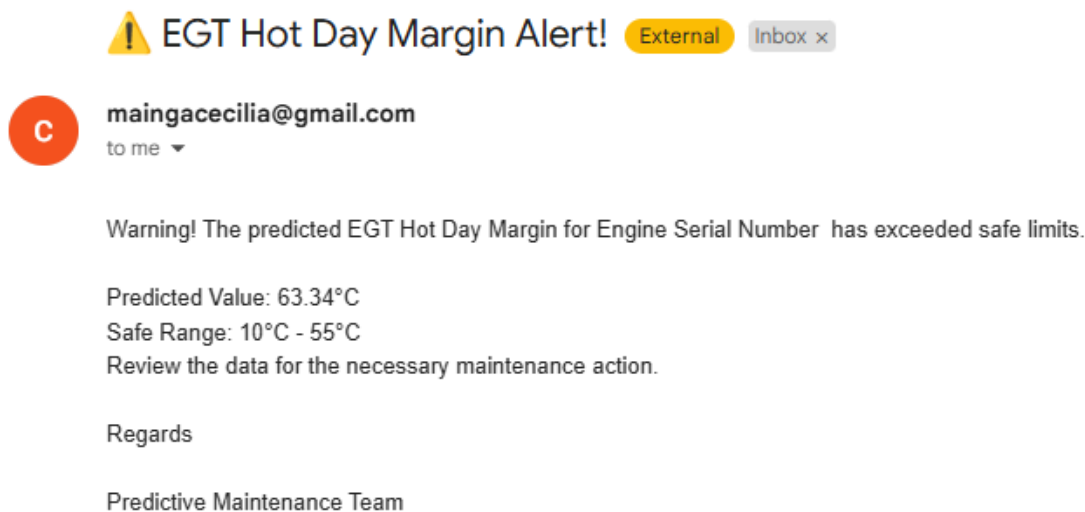


Figure 5.4: Email Notification

The deployment of the EGT prediction model using a Streamlit app is proof of concept to demonstrate the model's predictive capabilities in a user-friendly interface. Integration with onboard systems, however, would involve using APIs, or embedding the model into existing Aircraft Health Monitoring Systems. This would require close collaboration with system engineers to ensure compatibility with existing data pipelines.

5.2 Testing

This phase was conducted to validate the overall performance, usability, compatibility, security, and practical effectiveness of the EGT prediction application. The aim was to

ensure that the system addresses the intended problem of aircraft engine performance monitoring and meets the research objectives.

5.2.1 Functionality Testing

Functionality testing aimed to confirm that all components of the system operated correctly. Various samples of engine operational data were tested using both the file upload feature and the manual data entry option. In each case, the system successfully generated diverse and realistic EGT margin predictions, demonstrating reliable performance across input methods.

The email notification feature was also evaluated. After predictions were generated, users could input various valid email addresses to receive a summary containing the engine serial number, the threshold EGT value, and the predicted temperature. The results were properly formatted and successfully sent via email.

5.2.2 Usability Testing

The usability of the Streamlit web portal was assessed by navigating key features such as the input form, results display, SHAP-based interpretability charts, and the email notification option. The researcher evaluated how intuitive the application was for end-users, particularly non-technical staff in aviation contexts. Particular emphasis was placed on refining the wording of the instructions provided on the Overview page, which guides users on how to navigate and interact with the entire application. Clear placeholder text and step-by-step instructions were added to guide users.

5.2.3 Compatibility Testing

Compatibility testing ensured that the Streamlit-based application performed consistently across different browsers and devices. The application was tested on Google Chrome, Mozilla Firefox, and Microsoft Edge. Each browser successfully rendered the web portal, accepted input values, executed model predictions, and displayed SHAP plots without distortion or data loss. The prediction form, result visualization, and email functions all worked seamlessly across platforms.

5.2.4 Security Testing

Security measures were tested to ensure the safe handling of user interactions and sensitive operational data. Prior to accessing the application, users must log in using a username and password, providing an important layer of authentication to prevent unauthorized access. Although the application does not store personally identifiable information, the login system helps safeguard access to the prediction model and associated data.



Chapter 6: Discussion of Results

6.1 Explanatory Data Analysis Results

The box-plot in Figure 6.1 shows a box plot of the numerical columns of the data before scaling.

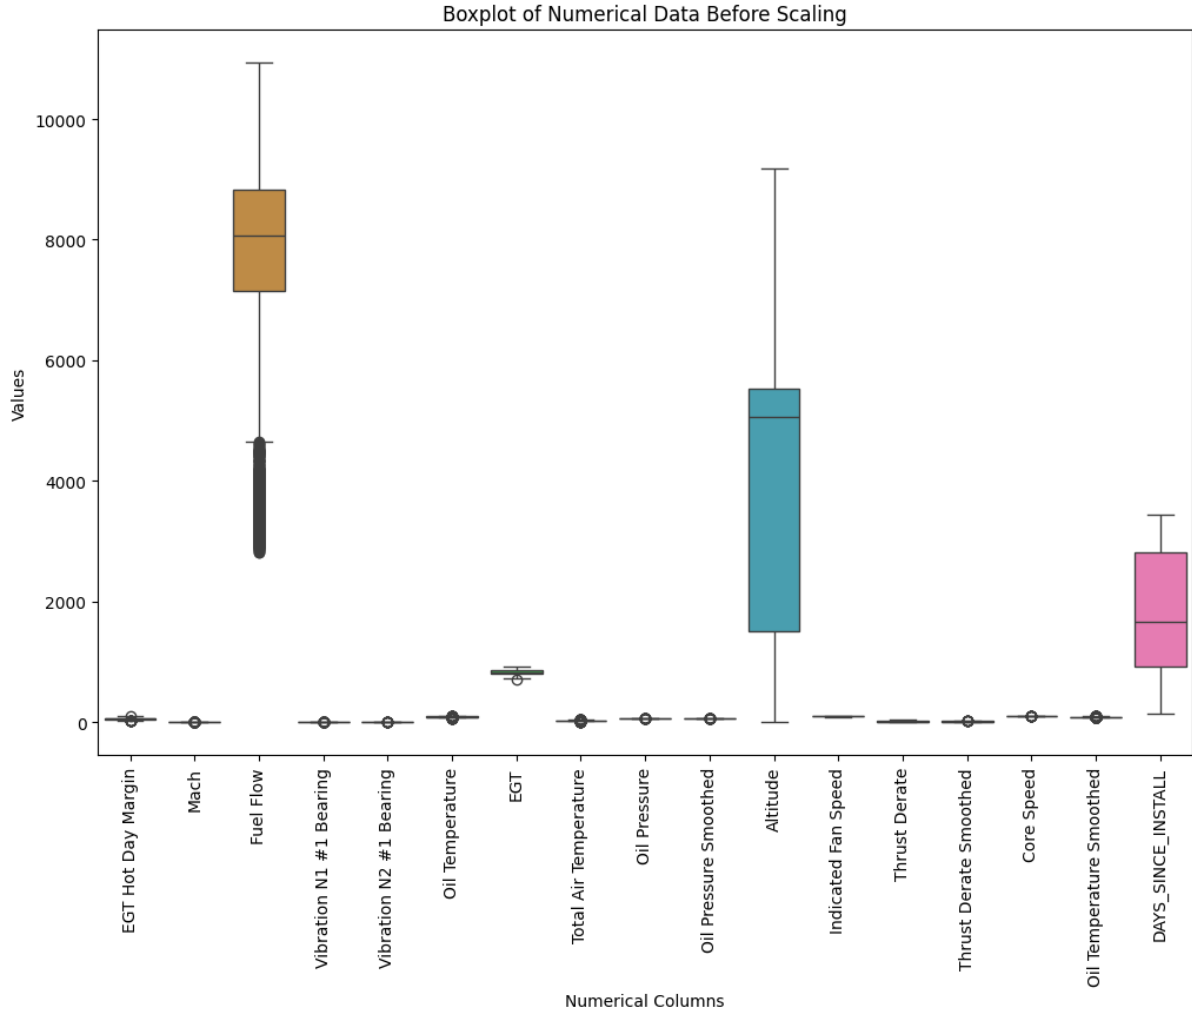


Figure 6.1: Boxplot of numerical columns before scaling

The boxplot did not provide clear insights due to the differing ranges of the variables, so min-max scaling was applied to enable more effective outlier detection.

The box-plot in Figure 6.2 below shows how outliers were detected after Min-Max scaling.

After scaling, outliers were detected in all the variables. The data used in this analysis represents actual engine data performance. Therefore outliers are a true reflection of the




Figure 6.2: Boxplot after min-max scaling

reality. Treating such outliers may cause the prediction model to be spurious.

The correlation heatmap in Figure 6.3 shows a strong negative(-0.83) relationship between Days-Since-Install and the the EGT Hot Day Margin. As the engine time increases, engine components such as compressor and turbine blades degrade and wear off. This reduces the overall component efficiency. To achieve the required thrust, the engine uses more power, gradually increasing the EGT and ultimately reducing the EGT Hot Day Margin. (FAA)

The high positive correlation (0.93) between Exhaust Gas Temperature (EGT) and indicated fan speed in an aircraft engine suggests that as the fan speed increases, the exhaust temperature also increases. The fan speed directs airflow at the engine intake. Higher fan speeds results in increased airflow into the engine and consequently, higher combustion temperatures as described by the Brayton cycle, which governs jet engine operation. This increased airflow and combustion leads to more energy being produced, which results in a higher exhaust temperature Cengel and Boles (2015). Additionally, there is a moderately negative (-0.68) correlation between Temperature and Altitude showing that ss altitude

increases, the temperature decreases. Sadraey (2017)




Figure 6.3: Correlation heatmap

6.2 Machine Learning Results

According to Table 6.1, deep learning model has the lowest values for MSE, RMSE, and MAE, as well as the highest R^2 score and Adjusted R^2 Score. This shows that the deep learning is the best-performing algorithm as it consistently shows the best performance across all metrics.

Table 6.1: Model performance metrics

Model	MAE	MSE	RMSE	R2	Adjusted R2
Random Forest	0.0541	0.005	0.070	0.766	0.764
Support Vector Regressor	0.033	0.002	0.042	0.913	0.913
K-Nearest Neighbors	0.046	0.004	0.060	0.825	0.823
Decision Tree	0.057	0.005	0.0735	0.741	0.739
Linear Regression	0.013	0.0003	0.0176	0.985	0.985
Deep Learning	0.0067	0.0001	0.0111	0.994	0.994

The Training and Validation loss plot for the original deep learning model in Figure 6.4. show that both the training and validation mean squared error decrease sharply during

the initial epochs, indicating that the model is learning effectively during this phase. After approximately 15 epochs, both metrics stabilize and fluctuate minimally, indicating convergence. The training and validation errors follow a similar trend, suggesting there is no significant overfitting or underfitting [Géron \(2022b\)](#).




Figure 6.4: Training and Validation loss plot for the original deep learning model

However, the loss plot of the optimized model in Figure 6.5 shows a gap between the curves. This means that the model performs significantly better on the training data than on the validation data, indicating overfitting [Géron \(2022b\)](#).

6.2.1 Class Imbalance

The examination of class imbalance in the dataset revealed no significant disproportion between risky instances (1) and non-risky instances (0), with 1320 class "1" counts representing approximately 66.73% of the dataset, and 658 class "0" counts accounting for approximately 33.27% of the dataset as shown in Figure 6.6. Therefore, no class balancing techniques were applied.

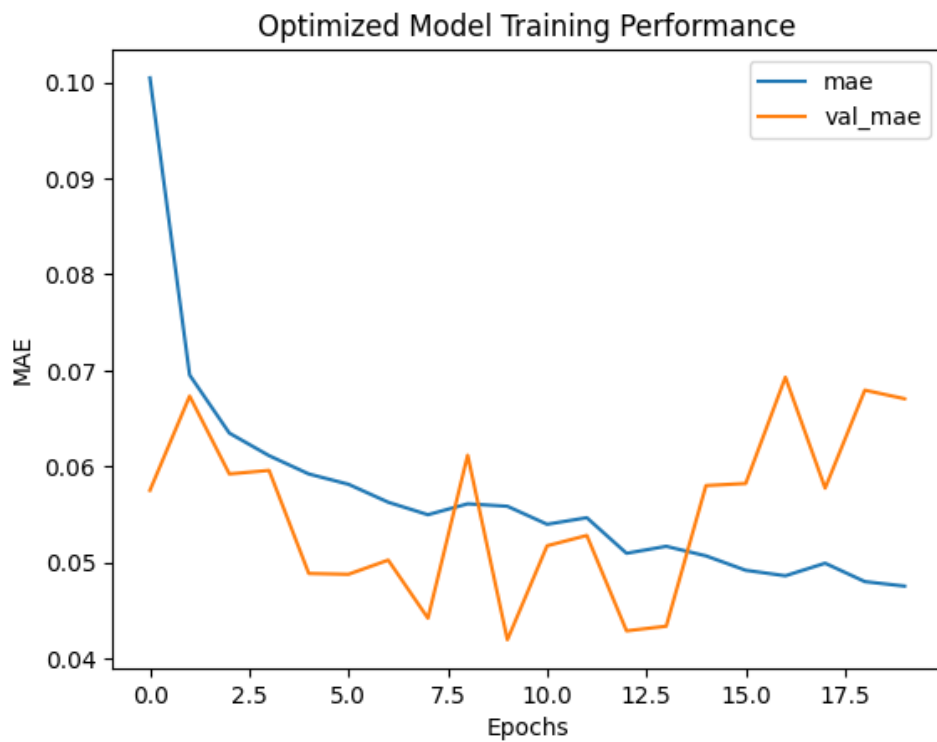


Figure 6.5: Training and validation loss plot of the optimized deep learning model

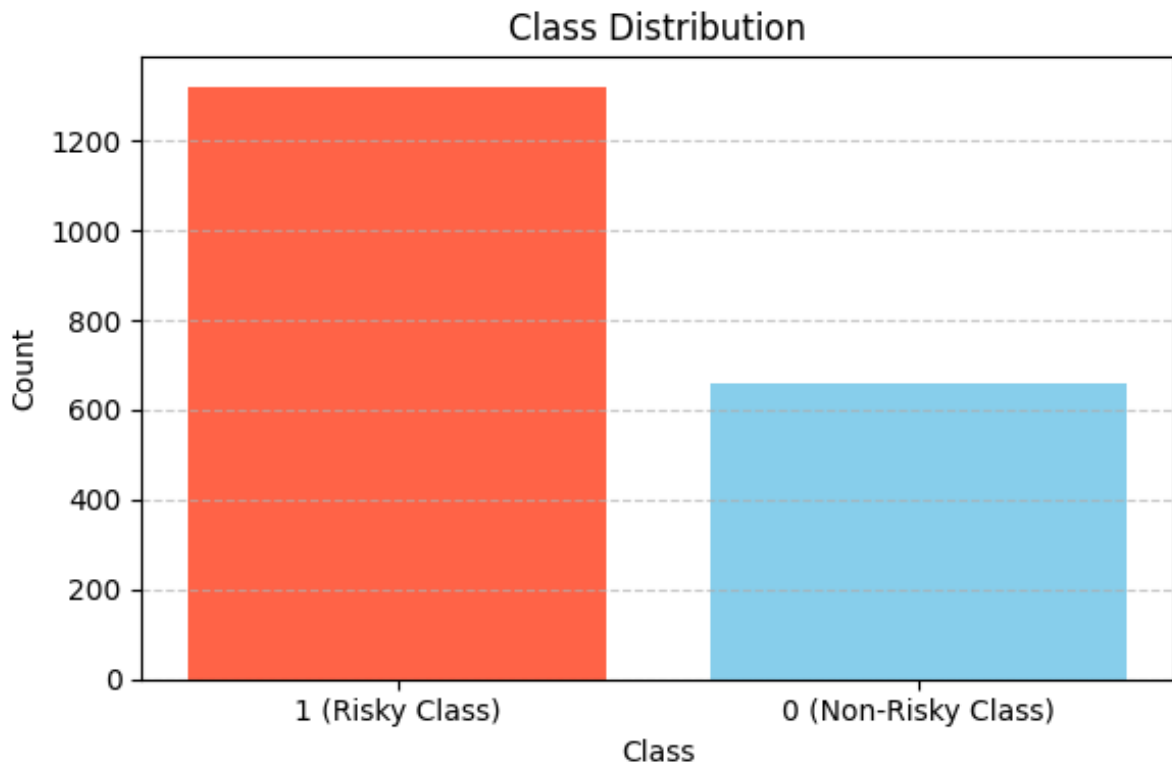


Figure 6.6: Class distribution plot

6.2.2 Comparison with Industry Limits

The confusion matrix of the original deep learning model in Figure 6.7, where class '1' represents the risky class and class '0' represents the non-risky class, shows that:

- (i) **True Positives (TP)**: The model accurately predicted 1 1315 times when the actual class was '1'.
- (ii) **False Positives (FP)**: The model wrongly classified five instances as '1' when they were '0'.
- (iii) **False Negatives (FN)**: The model mistakenly predicted '0' 26 times when the actual class was '1'.
- (iv) **True Negatives (TN)**: The model correctly predicted '0' 632 times when the actual class was '0'.

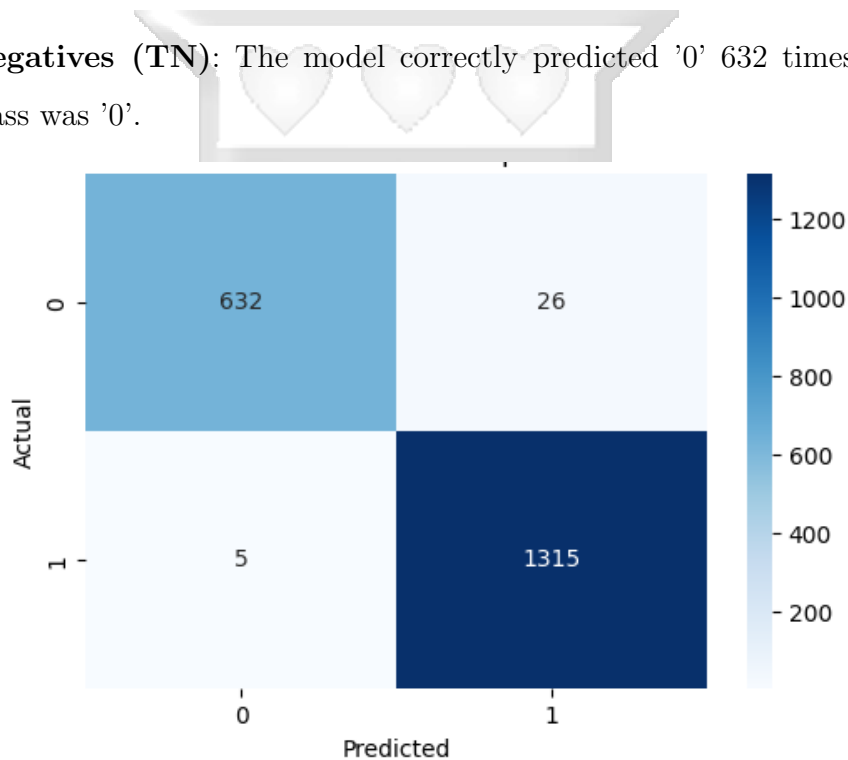


Figure 6.7: Confusion matrix of original deep learning model

This translates to an overall accuracy of 98% as shown in the classification report in Table 6.2. The report further indicates that the deep learning model achieved a high precision of 0.98 and a recall of 1.00 for the '1' class, and a precision of 0.99 for the '0' class. Similarly, the model has an F1-score of 0.99 and 0.98 for the '1' and '0' class respectively.

Table 6.2: Classification Report of the Original Deep Learning Model

Class	Precision	Recall	F1-Score	Support
0	0.99	0.96	0.98	658
1	0.98	1.00	0.99	1320
Accuracy	0.98			1978
Macro Avg	0.99	0.98	0.98	1978
Weighted Avg	0.98	0.98	0.98	1978

A comparison with the confusion matrix of the optimized deep learning model shown in Figure 6.8 reveals that the optimized model had a higher number of false positives and negatives, 94 and 54, respectively.




Figure 6.8: Confusion Matrix of the Optimized deep learning model

Furthermore, the classification report in Table 6.3 shows a lower precision of 0.87 in the '0' class as compared to the '1' class, where the precision is 0.96 and the recall is 0.93. Furthermore, there is a slight difference between the F1-scores of class '1' and class '0', with values of 0.94 and 0.89, respectively.

Table 6.3: Classification report of the optimized deep learning model

Class	Precision	Recall	F1-Score	Support
0	0.87	0.92	0.89	658
1	0.96	0.93	0.94	1320
Accuracy	0.93			1978
Macro Avg	0.91	0.92	0.92	1978
Weighted Avg	0.93	0.93	0.93	1978

The evaluation results demonstrate that the original deep learning model achieved ex-

cellent performance, particularly in detecting risky EGT margin cases. A key metric in this context is recall, which measures the model's ability to correctly identify all actual risky instances. The original model's recall of 1.00 for the risky class indicates that it successfully detected all cases where the EGT margin was truly risky, missing none. This is critical in operational settings because failing to detect a risky EGT margin could result in undetected engine issues, potentially leading to damage or safety hazards.

Moreover, the model showed a high precision of 0.98 for the risky class, meaning that when the model predicted a risky condition, it was correct 98% of the time. This reduces the likelihood of false alarms, which could otherwise lead to unnecessary maintenance or operational interruptions. The balance between high recall and high precision demonstrates the model's effectiveness in both minimizing missed detections and reducing false warnings.

6.3 SHAP Analysis Results

6.3.1 SHAP Summary Plot

The SHAP summary plots in Figure 6.9 display subplots that correspond to the machine learning models used in the project: Random Forest, Support Vector Regressor, KNN, Decision Tree, Linear Regression and Deep Learning. For each model, the y-axis lists the features used in the model while the x-axis represents the SHAP value, which indicates how much a given feature contributed to a specific prediction. Each dot represents a single data point, with the color representing feature values: red represent high values while blue represent low values.

The key features: days-since-install, EGT, oil temperature, Indicated Fan Speed, Total Air Temperature, altitude and thrust derate influence SHAP values as described below:

- (i) Days-Since-Install: When engine time increases, the SHAP value decrease hence decreasing the EGT Hot Day Margin, and vice versa. The inverse relationship between engine time and SHAP value, leading to a decrease in the EGT Hot Day Margin, can be explained by engine wear and performance degradation over time. as corroborated by the correlation heatmap in Figure 4.3. Due to wear, it operates less efficiently and requires more effort to maintain performance, resulting in a gradual increase in EGT over time. This, in turn, causes a reduction in the EGT

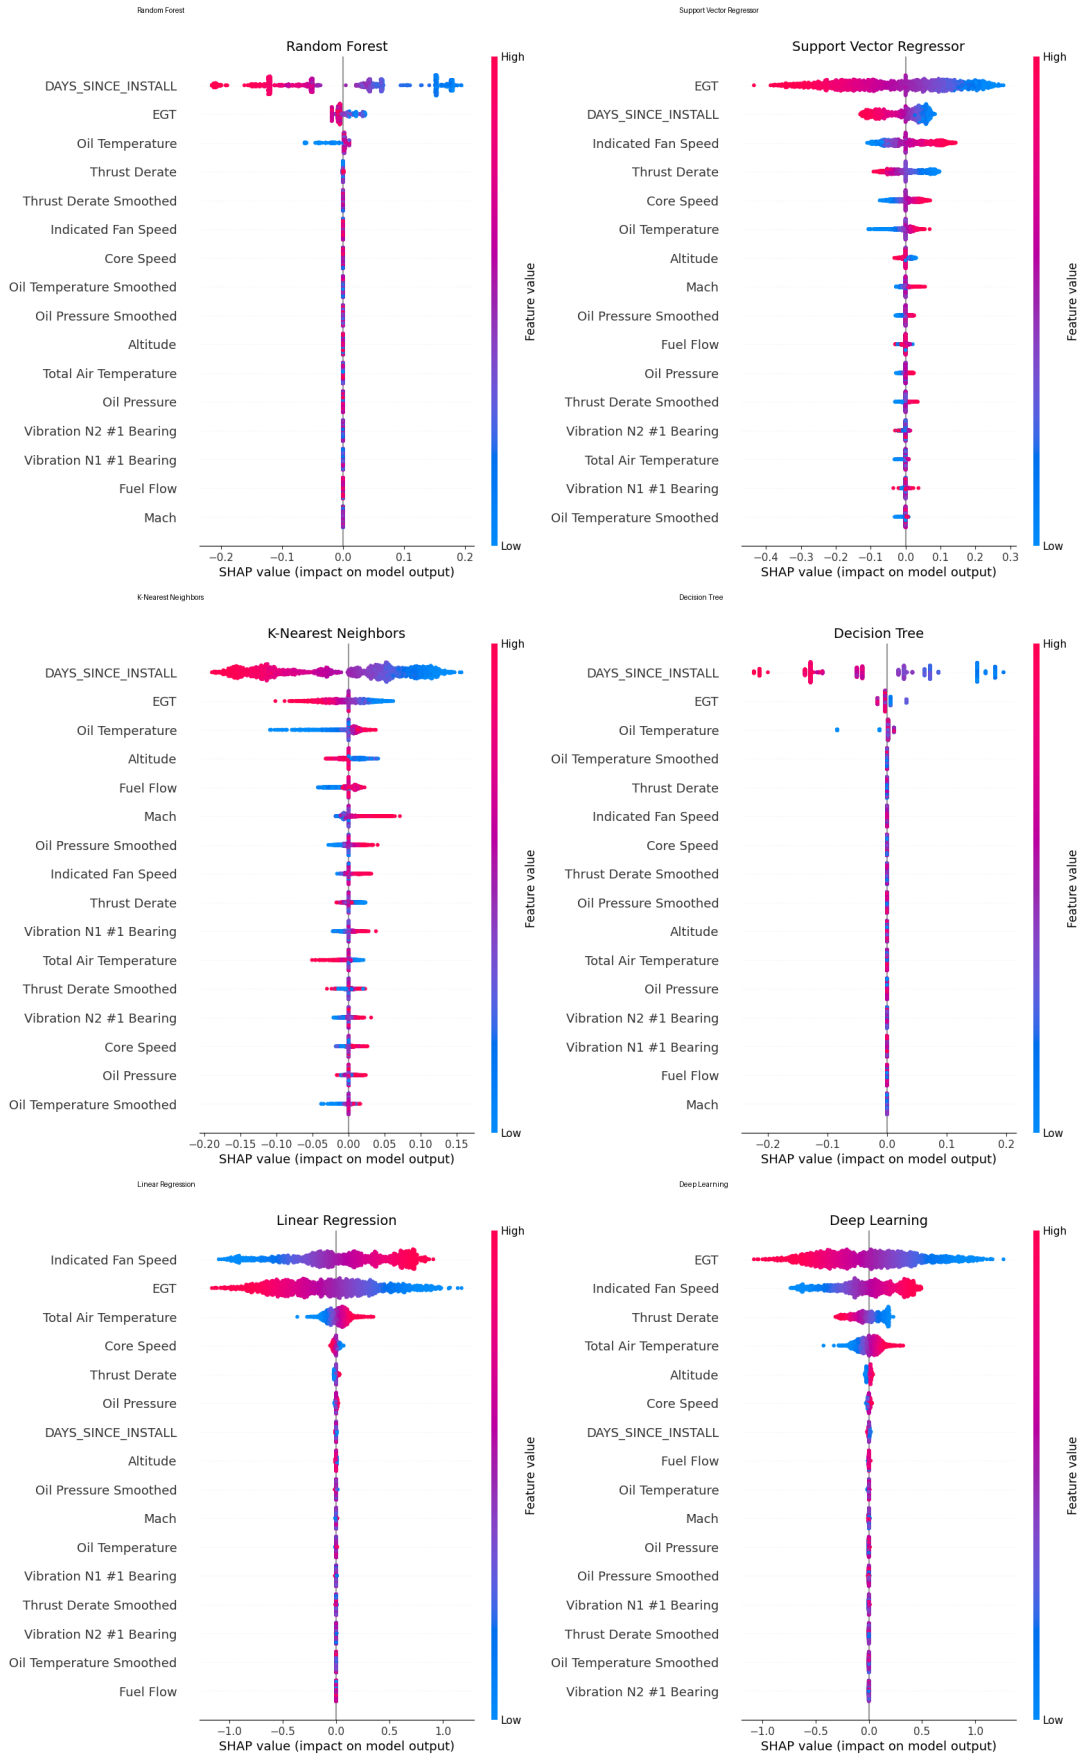


Figure 6.9: SHAP summary plot

margin (FAA).

- (ii) EGT : As the EGT rises, the SHAP value declines, leading to a decrease in the EGT Hot Day Margin. This inverse relationship is due to the definition of EGT margin, which is the difference between the engine’s operating EGT and the maximum allowable limit set by the manufacturer. Consequently, a higher EGT results in a lower EGT margin [Yildirim and Kurt \(2016\)](#).
- (iii) Indicated Fan speed: As the fan speed increase, SHAP values also increase and vice versa. This relationship aligns with the principles of the Brayton cycle, which governs jet engine operation where improving the efficiency of the Bryton cycle can be achieved by Increasing the turbine inlet (or firing) temperatures. Higher efficiency leads to greater thrust, resulting in increased EGT and consequently a lower EGT margin [Cengel and Boles \(2015\)](#).
- (iv) Total Air Temperature: In the Support Vector Regression, Deep Learning, and Linear Regression models, an increase in air temperature corresponds to higher SHAP values, indicating a rise in EGT Hot Day Margin. This contradicts thermodynamic principles, which state that higher outside air temperature leads to higher EGT, thereby reducing the EGT margin. The observed relationship in these models suggests the presence of operational adjustments, such as thrust derate, which may influence the EGT margin. However, the K-Nearest Neighbors model correctly captures the expected trend—an increase in OAT leads to lower SHAP values, indicating a reduced EGT margin. [Balicki et al. \(2014\)](#).
- (v) Thrust Derate: As thrust derate decreased, SHAP values increased. Thrust derate calculated as a percentage reduction from the maximum rated thrust (FAA), hence increasing the derate reduces the thrust level. Reduced thrust results in EGT, thereby increasing the EGT margin.

6.3.2 SHAP Feature Importance plot

Figure 6.10 shows SHAP Feature importance plots for each model, illustrating the mean absolute SHAP values for different features. The dominant features across the model are days-since-install and EGT, highlighting their significant impact on EGT Hot Day Margin predictions. Other influential features include Indicated Fan Speed, which emerged as

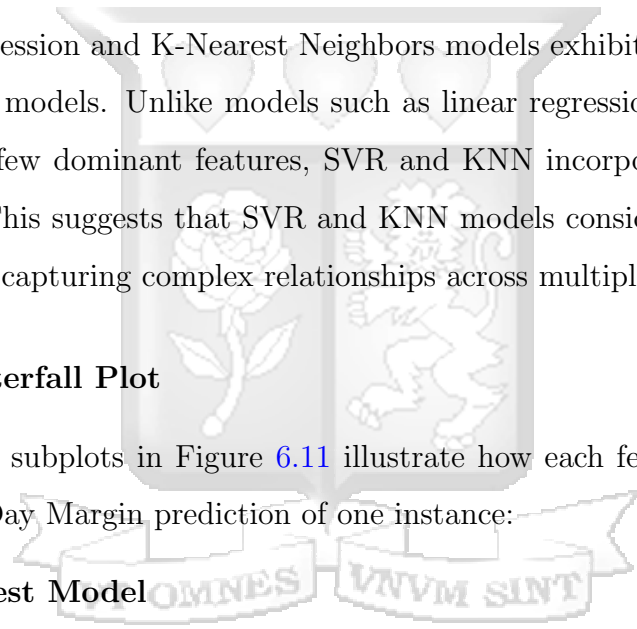
a key factor, particularly in K-nearest neighbors, Support Vector Regression, and Deep Learning. Similarly, Total Air Temperature contributes notably, especially in Linear Regression and Deep Learning, reinforcing the importance of ambient conditions in engine performance.

Thrust Derate is another critical factor, which appears in Support Vector Regression and K-Nearest Neighbor models. Its significance suggests that operational adjustments, such as reducing thrust, modify the EGT margin. Similarly, oil temperature appears to be an influential factor across models, particularly in KNN and decision trees, suggesting its relevance in monitoring engine thermal behavior. Lastly, Core Speed shows moderate influence, particularly in SVR and Decision Tree models.

Support Vector Regression and K-Nearest Neighbors models exhibited a broader feature influence than other models. Unlike models such as linear regression or Decision Trees, which focused on a few dominant features, SVR and KNN incorporated a wider range of input variables. This suggests that SVR and KNN models consider a holistic view of engine performance, capturing complex relationships across multiple parameters.

6.3.3 SHAP Waterfall Plot

The SHAP waterfall subplots in Figure 6.11 illustrate how each feature contributed to the final EGT Hot Day Margin prediction of one instance:

- 
- a) **Random Forest Model**
 - (i) Oil Temperature (-0.04): Higher oil temperatures decreased the predicted EGT Hot Day margin.
 - (ii) Days-Since-Install (+0.04): Longer engine operational days increased the predicted EGT Hot Margin.
 - (iii) EGT (-0.01): Higher EGT values lower the predicted EGT Hot Day margin.
 - b) **Support Vector Regressor Model**
 - (i) EGT (-0.19): Higher EGT reduced the EGT Hot Day margin predicted.
 - (ii) Indicated Fan Speed (+0.09): Higher fan speed increased the predicted margin.




Figure 6.10: SHAP feature importance plot

SHAP Waterfall Plots for All Models (First Instance)




Figure 6.11: SHAP Waterfall plot

- (iii) Thrust Derate (+0.07): More derate increased EGT Hot Day Margin predicted.
- (iv) Oil Temperature (-0.06): Higher oil temperature lowered the predicted margin.
- (v) Core Speed (-0.06): Increased core speed reduced EGT margin.
- (vi) Mach (+0.02): Higher Mach values increased the predicted margin slightly.

c) **K-Nearest Neighbor Model**

- (i) Oil Temperature (+0.07): Higher oil temperatures increased predicted EGT Hot Day margin.
- (ii) Days-Since-Install (+0.04): Longer engine operational periods led to higher predicted EGT Hot Day margin.
- (iii) EGT (-0.03): Higher EGT reduced the predicted margin.
- (iv) Mach (+0.03): Higher Mach increased EGT margin.

d) **Decision Tree Model**

- (i) Days-Since-Install (-0.04): Longer engine operational periods led to higher predicted EGT Hot Day margin.
- (ii) Oil Temperature (-0.01): Higher oil temperature slightly decreased the predicted EGT Hot Day margin.

e) **Linear Regression Model**

- (i) Indicated Fan Speed (N1) (+0.69): Higher fan speed significantly increased the predicted EGT Hot Day margin.
- (ii) Total Air Temperature (-0.59): Higher ambient temperature reduced the predicted EGT Hot Day margin.
- (iii) Core Speed (-0.02): Increased core speed reduced the predicted EGT margin.
- (iv) Thrust Derate (-0.02): More derate decreased the predicted EGT margin.

f) **Deep Learning Model**

- (i) EGT (-0.49): Higher EGT reduced the predicted margin.

- (ii) Indicated Fan Speed (+0.38): Higher fan speed increased the predicted margin.
- (iii) Thrust Derate (+0.18): High derate increased the predicted EGT Hot Day margin.
- (iv) Core Speed (-0.1): Increased core reduced the predicted margin slightly.
- (v) Mach (+0.03): Higher Mach increased the predicted margin.

Across all models, EGT consistently decreased the predicted EGT margin, while Indicated Fan Speed increased it. Other features, such as thrust derate, mach oil temperature, thrust derate, and Days-Since-Install show varying effects across models. This detailed breakdown of how individual features influence a model's prediction for a single instance are also illustrated in Figures 6.12 to 6.16.




Figure 6.12: Random forest force plot




Figure 6.13: Decision tree force plot



Figure 6.14: KNN force plot




Figure 6.15: Linear regression force plot




Figure 6.16: Support vector regressor force plot

The decision subplots in Figure 6.17 show how the top 10 features influenced every model's decision path for a single prediction instance. All models except K-nearest Neighbor and Support Vector Regressor made decisions based on the dominant features; EGT, Total Air Temperature (TAT), Oil temperature and Oil pressure. However, K-Nearest Neighbors (KNN) and Support Vector Machine (SVM) incorporated all the features in the decision path.

SHAP Decision Plots for All Models(Top 10 Features)

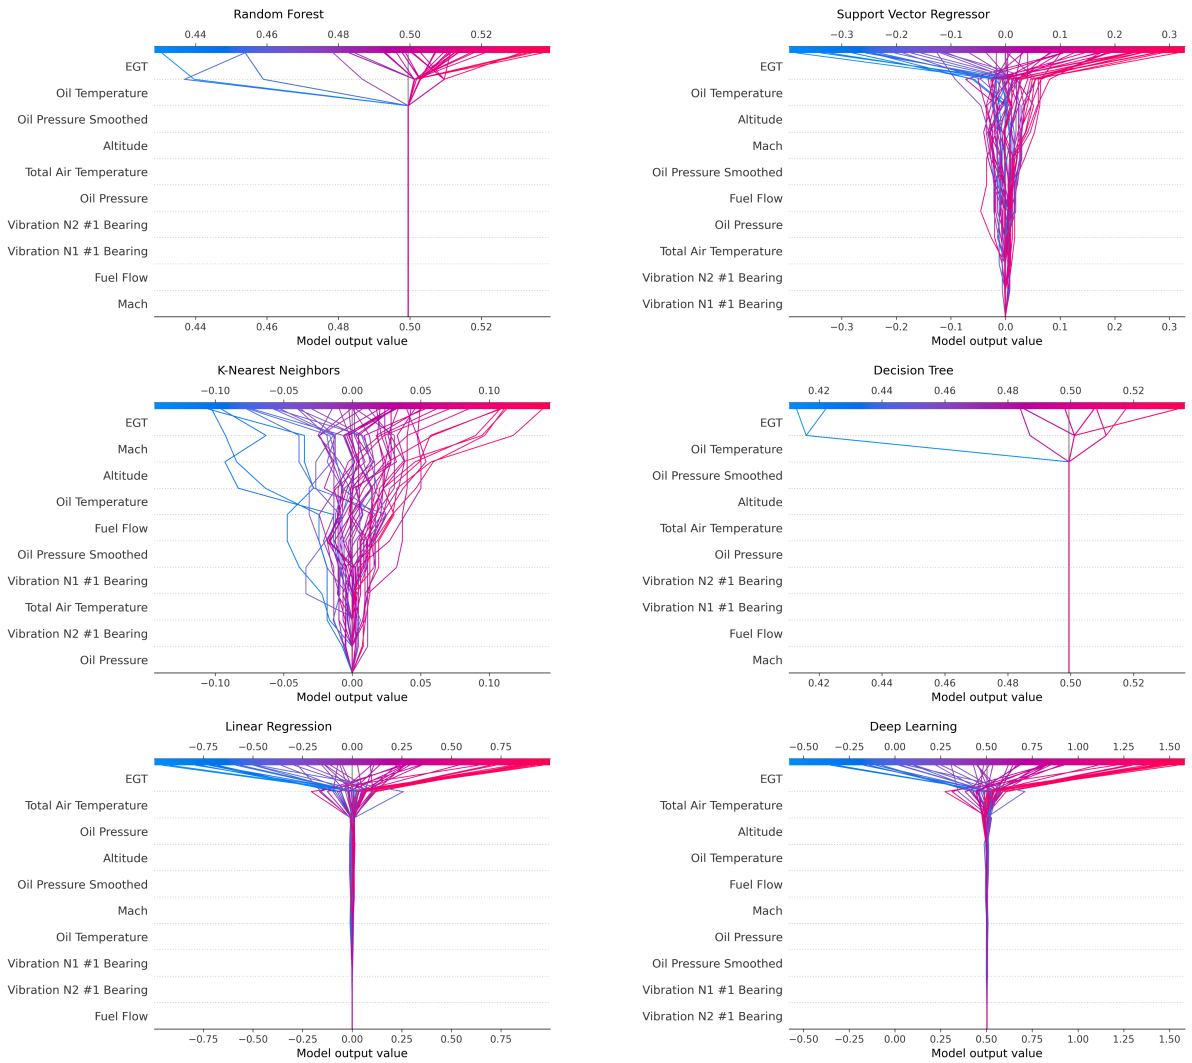


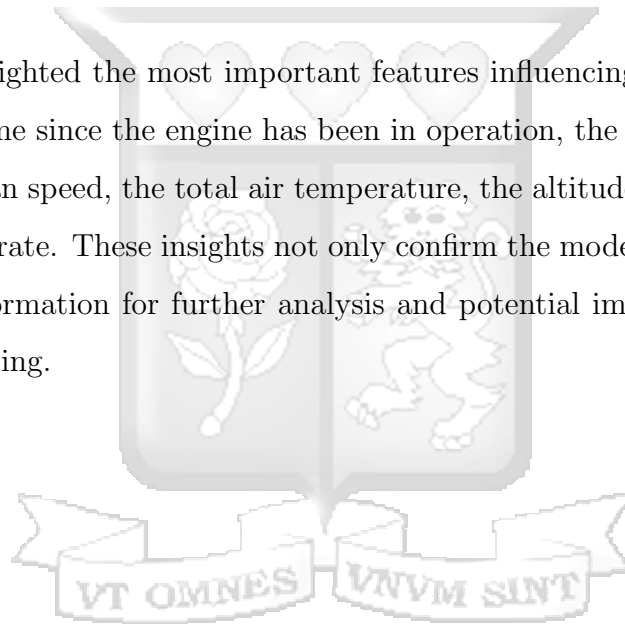
Figure 6.17: SHAP decision plots

6.4 Summary

In summary, the deep learning model was the best performing model. It achieved the lowest values for mean absolute error, mean squared error, and root mean squared error, while also attaining the highest scores for the coefficient of determination and the adjusted coefficient of determination. These results indicate that the model provided the most accurate predictions and the best overall fit to the data.

Efforts to improve the model further through hyperparameter tuning did not result in any significant enhancement of its performance. As a result, the original version of the deep learning model was selected for deployment on the Streamlit Community Cloud platform.

SHAP analysis highlighted the most important features influencing the model's output. These include the time since the engine has been in operation, the exhaust gas temperature, the indicated fan speed, the total air temperature, the altitude, the core speed, and the level of thrust derate. These insights not only confirm the model's reasoning but also provide valuable information for further analysis and potential improvements in engine performance monitoring.



Chapter 7: Conclusions and and Future Work

7.1 Conclusion

This study explored the application of predictive maintenance techniques to aircraft engines, specifically focusing on the CFM56-7B26E engines used in the Boeing 737NG aircraft. By analyzing historical engine performance data, several regression models were used to predict the Exhaust Gas Temperature Hot Day Margin, a critical indicator of engine health and structural integrity. Regression models, including Linear Regression, Decision Tree, K-Nearest Neighbors, Random Forest, Support Vector Machine, and Neural Networks, were trained on the data, and their performance was evaluated.

The analysis revealed that the deep neural network performed better than other models, in terms of mean absolute error, mean square error and root mean square error. Hence, the deep learning model was used as a backend to deploy a prediction app on Streamlit community cloud platform. The web app allows aviation professionals to improve operational efficiency, enhance flight safety and planning, and optimize maintenance planning by combining predictive analytics, an intuitive user interface, and alerts.

Furthermore, SHAP feature importance analysis revealed that EGT, Days-since-install (how long the engine has been operational), Indicated Fan Speed, Oil Temperature, and Total Air Temperature are the top five most important features. These features play a critical role in influencing the models' predictions, highlighting their significance in assessing engine health and performance.

This project highlights the potential of applying machine learning in the aviation maintenance industry for predictive maintenance, driving safer, more efficient, and cost-effective operations. Early detection of abnormal engine behavior not only reduces the risk of unexpected failures but also minimizes maintenance costs and extends the lifespan of critical engine components. By implementing such predictive models, airlines can optimize their maintenance schedules, enhancing overall fleet reliability and reducing downtime.

7.2 Future Work

Several avenues for future work are proposed to enhance the robustness and applicability of the EGT margin prediction model. First, the current model was developed using a

dataset limited to a single engine type, which may restrict its generalization across different engine platforms. Future efforts should focus on expanding the dataset to include multiple engine models to improve the model's adaptability and predictive performance across varied operating conditions. Secondly, while the model demonstrates strong predictive capabilities, it has not yet been validated in a real-time environment. Integrating real-time data streaming and testing the model within a live engine health monitoring framework will be crucial for operational deployment. Improving the interpretability of the model's outputs, particularly for non-technical stakeholders such as maintenance engineers and operations personnel, would increase its usability and support informed decision-making in safety-critical scenarios.



Bibliography

- (2008). *Owner's Operator's Guide:CFM56-7B*.
- (2012). *Aviation Maintenance Technician Handbook-powerplant*. Aviation Supplies & Academics.
- Ackert, S. (2011). Engine maintenance concepts for financiers-elements of turbofan shop maintenance costs (tech. rep.).
- Ackert, S. P. (2010). Basics of aircraft maintenance programs for financiers. *Evaluation & insights of commercial aircraft maintenance programs*, 10.
- Adryan, F. and Sastra, K. (2021). Predictive maintenance for aircraft engine using machine learning: trends and challenges. *Avia*, 3(1).
- Ahmed, M., Badshah, S., Khan, R. U., Sajjad, M., Jan, S., and Pakistan, I. (2015). Design and simulation of flow through subsonic axial compressor.
- Apostolidis, A., Bouriquet, N., and Stamoulis, K. P. (2022). Ai-based exhaust gas temperature prediction for trustworthy safety-critical applications. *Aerospace*, 9(11).
- Association, I. A. T. (2022). Financial performance of the aviation industry 2022. Accessed: 2024-09-24.
- Badea, V. E., Zamfiroiu, A., and Boncea, R. (2018). Big data in the aerospace industry. *Informatica Economica*, 22(1):17–24.
- Balicki, W., Głowacki, P., Szczecinski, S., Chachurski, R., and Szczeciński, J. (2014). Effect of the atmosphere on the performances of aviation turbine engines. *Acta Mechanica et Automatica*, 8:70–73.
- Boyce, M. P. (2011). *Gas turbine engineering handbook*. Elsevier.
- Cai, J., Li, X., and Chen, X. (2016). Optimization of aeroengine shop visit decisions based on remaining useful life and stochastic repair time. *Mathematical Problems in Engineering*, 2016(1):5648360.
- Cengel, Y. A. and Boles, M. A. (2015). *Thermodynamics: An Engineering Approach*. McGraw-Hill Education.

- Chicco, D., Warrens, M. J., and Jurman, G. (2021). The coefficient of determination r-squared is more informative than smape, mae, mape, mse and rmse in regression analysis evaluation. *Peerj computer science*, 7:e623.
- Daggett, D. L., Diedrich, M., and Duong, H. (2003). Engine technology for reducing aircraft fuel consumption. Technical report, NASA.
- Ehrich, F. F. and Baxter, . A. D. (2024). britannica.com. Accessed: 2024-09-10.
- Emer, N. and Özbek, N. (2020). A survey on kalman filtering for unmanned aerial vehicles: Recent trends, applications, and challenges. In *Proceedings of the International Conference on Engineering Technologies (ICENTE'20), Konya, Turkey*, pages 19–21.
- (FAA), F. A. A. (1988). Reduced and derated advisory circular. Retrieved 2025.
- (FAA), F. A. A. (2012). *Aviation Maintenance Technician Handbook - Powerplant Volume 2. Aviation Supplies & Academics*.
- Fabry, S. and Češkovič, M. (2017). Aircraft gas turbine engine vibration diagnostics. *MAD-Magazine of Aviation Development*, 5(4):24–28.
- Fedele, L., Di Vito, L., and Ramundo, F. E. (2020). Increasing efficiency in an aeronautical engine through maintenance evaluation and upgrades: Analysis of the reliability and performance improvements under financial issues. *Energies*, 13(12):3059.
- Géron, A. (2022a). *Hands-on Machine Learning with Scikit-Learn, Keras, and TensorFlow*. "O'Reilly Media, Inc."
- Géron, A. (2022b). *Hands-on Machine Learning with Scikit-Learn, Keras, and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems*. "O'Reilly Media, Inc."
- Hackeling, G. (2017). *Mastering Machine Learning with scikit-learn*. Packt Publishing Ltd.
- Hu, H., Shyy, W., and Shih, T. I. (2010). Lift, thrust, and flight. *Encyclopedia of Aerospace Engineering*, 2:837.

- International Civil Aviation Organization (2022). The world of air transport in 2022. Accessed: 2024-09-25.
- International Civil Aviation Organization (ICAO) (2023). Global aviation safety plan (gasp). Accessed: 2024-09-26.
- Kefalas, M., Santiago Rojo Jr, J. D., Apostolidis, A., Van Den Herik, D., van Stein, B., and Bäck, T. (2022). Explainable artificial intelligence for exhaust gas temperature of turbofan engines. *Journal of Aerospace Information Systems*, 19(6):447–454.
- Keivanpour, S. and Kadi, D. A. (2015). A sustainable approach to aircraft engine maintenance. *IFAC-PapersOnLine*, 48(3):977–982.
- Kenya Civil Aviation Authority (2024). Passenger traffic statistics. <https://www.kcaa.or.ke/about-us/statistics/passenger-traffic>. Accessed: 2024-10-01.
- Koc, H., Erdoğan, A., Barjakly, Y., and Peker, S. (2021). Uml diagrams in software engineering research: A systematic literature review. *Proceedings*, 74:13.
- Kurz, R. and Brun, K. (2001). Degradation in gas turbine systems. *J. Eng. Gas Turbines Power*, 123(1):70–77.
- Lee, J. and Mitici, M. (2023). Deep reinforcement learning for predictive aircraft maintenance using probabilistic remaining-useful-life prognostics. *Reliability Engineering & System Safety*, 230:108908.
- Liu, S., Zhou, N., Song, C., Chen, G., and Wu, Y. (2024). Exhaust gas temperature prediction of aero-engine via enhanced scale-aware efficient transformer. *Aerospace*, 11(2).
- Lundberg, S. M. and Lee, S.-I. (2017). A unified approach to interpreting model predictions. *Advances in neural information processing systems*, 30.
- Ma, S., Wu, Y., Zheng, H., and Gou, L. (2023). A hybrid of narx and moving average structures for exhaust gas temperature prediction of gas turbine engines. *Aerospace*, 10(6).

- Matthaiou, I., Khandelwal, B., and Antoniadou, I. (2017). Vibration monitoring of gas turbine engines: Machine-learning approaches and their challenges. *Frontiers in Built Environment*, 3:54.
- Mattingly, J. D. (2002). *Aircraft engine design*. Aiaa.
- Mofokeng, T., Mativenga, P. T., and Marnewick, A. (2020). Analysis of aircraft maintenance processes and cost. *Procedia CIRP*, 90:467–472. 27th CIRP Life Cycle Engineering Conference (LCE2020) Advancing Life Cycle Engineering : from technological eco-efficiency to technology that supports a world that meets the development goals and the absolute sustainability.
- NASA Glenn Research Center (2022). Four forces on an airplane.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Reilly, T. (1996). The st. petersburg-tampa airboat line: 90 days that changed the world of aviation. *Tampa Bay History*, 18(2):4.
- Sadraey, M. H. (2017). *Aircraft performance: an engineering approach*. CRC Press.
- Sahin, S. (May 10th 2024). Scalers: Standard vs minmax. Accessed: 2024-07-16.
- Sehra, A. K. and Whitlow, W. (2004). Propulsion and power for 21st century aviation. *Progress in Aerospace Sciences*, 40(4):199–235.
- Sree, K. (April 2nd 2024). Missing data: Handling missing values in pandas with python. Accessed: 2024-07-16.
- Stanton, I., Munir, K., Ikram, A., and El-Bakry, M. (2023). Predictive maintenance analytics and implementation for aircraft: Challenges and opportunities. *Systems Engineering*, 26(2):216–237.
- Tatachar, A. V. (2021). Comparative assessment of regression models based on model evaluation metrics. *International Journal of Innovative Technology and Exploring Engineering*, 8(9):853–860.


Ullah, S., Li, S., Khan, K., Khan, S., Khan, I., and Eldin, S. M. (2023). An investigation of exhaust gas temperature of aircraft engine using lstm. *IEEE Access*, 11:5168–5177.

United Nations (2023). Goal 9: Build resilient infrastructure, promote inclusive and sustainable industrialization and foster innovation.

Wang, Z. and Zhao, Y. (2023). Data-driven exhaust gas temperature baseline predictions for aeroengine based on machine learning algorithms. *Aerospace*, 10(1).

Yildirim, M. T. and Kurt, B. (2016). Engine health monitoring in an aircraft by using levenberg-marquardt feedforward neural network and radial basis function network. In *2016 International Symposium on INnovations in Intelligent SysTems and Applications (INISTA)*, pages 1–5. IEEE.

Yildirim, M. T. and Kurt, B. (2018). Aircraft gas turbine engine health monitoring system by real flight data. *International Journal of Aerospace Engineering*, 2018(1):9570873.



Appendices

Appendix A: Similarity Report

168951- Predictive Maintenance of Aircraft Engines Using Machine Learning A Case of the CFM567B26E Engine.pdf

ORIGINALITY REPORT

15%

SIMILARITY INDEX

13%

INTERNET SOURCES

15%

PUBLICATIONS

10%

STUDENT PAPERS

PRIMARY SOURCES

1	su-plus.strathmore.edu Internet Source	2%
2	Submitted to Strathmore University Student Paper	1%
3	www.mdpi.com Internet Source	1%
4	dspace.lib.ntua.gr Internet Source	1%
5	Dasheng Xiao, Hong Xiao, Rui Li, Zhanxue Wang. "Application of physical-structure-driven deep learning and compensation methods in aircraft engine health management", Engineering Applications of Artificial Intelligence, 2024 Publication	<1%
6	transp-or.epfl.ch Internet Source	<1%
7	Ghaffarigharehbagh, Akram. "Studying the Mechanics of TiB via Molecular Dynamics and Machine Learning.", The University of Iowa Publication	<1%
8	commons.erau.edu Internet Source	<1%
9	Ribeiro, Francisco. "Machine Learning Applied to Deep Space Images", Universidade do	<1%

11	Internet	medium.com	<1%
12	Publication	Arvind Dagur, Karan Singh, Pawan Singh Mehra, Dharendra Kumar Shukla. "Intelli...	<1%
13	Submitted works	Heriot-Watt University on 2023-05-03	<1%
14	Internet	www.e-du.ru	<1%
15	Submitted works	University of North Texas on 2024-05-06	<1%
16	Publication	Makou, Maryam Bayat. "Search for Higgs Boson Production in Association with B-...	<1%
17	Publication	Shafi Ullah, ShuGuang Li, Khalid Khan, Shahbaz Khan, Ilyas Khan, Sayed M Eldin. "...	<1%
18	Submitted works	Tilburg University on 2025-05-15	<1%
19	Internet	hdl.handle.net	<1%
20	Internet	www.frontiersin.org	<1%
21	Submitted works	Edith Cowan University on 2023-05-30	<1%
22	Submitted works	Middle East Technical University on 2013-09-09	<1%
23	Submitted works	University College London on 2023-08-20	<1%
24	Submitted works	Cranfield University on 2013-09-05	<1%

Appendix B: Ethical Clearance Confirmation



17th January 2025

Ms Oketch Cecilia,
cecilia.oketch@strathmore.edu

Dear Ms Oketch,

RE: Predictive Maintenance of Aircraft Engines Using Machine Learning: A Case of the CFM56-7B26E Engine

This is to inform you that SU-ISERC has reviewed and approved your above SU-masters proposal. Your application reference number is SU-ISERC2486/24. The approval period is from 17th January 2025 to 16th January 2026.

This approval is subject to compliance with the following requirements:

- i. Only approved documents including (informed consents, study instruments, MTA) will be used.
- ii. All changes including (amendments, deviations, and violations) are submitted for review and approval by SU-ISERC.
- iii. Death and life-threatening problems and serious adverse events or unexpected adverse events whether related or unrelated to the study must be reported to SU-ISERC within 72 hours of notification.
- iv. Any changes anticipated or otherwise that may increase the risks or affected safety or welfare of study participants and others or affect the integrity of the research must be reported to SU-ISERC within 72 hours.
- v. Clearance for the export of biological specimens must be obtained from relevant institutions.
- vi. Submission of a request for renewal of approval at least 60 days prior to the expiry of the approval period. Attach a comprehensive progress report to support the renewal.
- vii. Submission of an executive summary report within 90 days of completion of the study to SU-ISERC.

Before commencing your study, you will be expected to obtain a research license from National Commission for Science, Technology, and Innovation (NACOSTI) <https://research-portal.nacosti.go.ke/> and obtain other clearances needed.

Yours sincerely,

Mr Ambrose Rachier,
Chairperson; SU-ISERC