Gas Path Analysis on the GEnx-1B Engine with Fewer Gas Path Sensors

MSc. Thesis at KLM Engine Services

Lucas Middendorp

TU Delft Delft University of Technology
KLM Engineering

Gas Path Analysis on the GEnx-1B Engine with Fewer Gas Path Sensors

MSc. Thesis at KLM Engine Services

by

Lucas Middendorp

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Friday June 12, 2026 at 10:00 AM.

Student number:	5067049	
Project duration:	August 2025 – June 2026	
Thesis committee:	Prof. dr. ir. P. Colonna,	TU Delft, chair
	Dr. ir. W.P.J. Visser,	TU Delft, supervisor
	Dr. ir. I. de Pater,	TU Delft
	ir. T.O. Rootliep	KLM, supervisor
	ir. D. Cisneros Acevedo	TU Delft

An electronic version of this thesis is available at `http://repository.tudelft.nl/`.

Executive Summary

The aviation industry is increasingly driven by the need to reduce maintenance costs and improve reliability. Achieving these objectives requires condition monitoring techniques that can extract reliable information on engine health. This master thesis investigates health indicator (HI) selection and prediction for the General Electric GEnx-1B turbofan engine, where health indicators represent deviations in component efficiency and mass flow capacity. The work builds upon Gas Path Analysis (GPA) methods and introduces a novel health indicator selection framework, in combination with the analysis of companion engines, to improve diagnostic performance.

Gas turbines operate under a wide range of operating conditions, leading to performance deterioration over time. Component deterioration is quantified using the HIs of four components in the GEnx-1B model: fan core + low-pressure compressor (LPC), high-pressure compressor (HPC), high-pressure turbine (HPT), and low-pressure turbine (LPT). A key challenge in condition monitoring of the GEnx-1B is the reduction of available gas path sensors. This leads to an underdetermined problem for GPA, where the number of unknown health indicators (8) exceeds the number of known measurements (6).

To address this issue, this thesis proposes a method called the Component Exclusion Method (CEM). This method combines the health indicator predictions of multiple subsets using a weight-factor matrix. A subset is defined as a combination of six out of eight health indicators, ensuring a solvable problem for each subset. By evaluating multiple subsets and weighting them based on their predictive accuracy, this method enables accurate prediction of health indicators. The modeling tool for the gas path analysis is GSPy, a Python implementation of the Gas turbine Simulation Program GSP. This tool has allowed for customizations and extensions to be introduced, which improved the model matching with the on-wing data. These include simulated deterioration, N1 control for adaptive modeling, crossflow implementation and measurement tolerances.

Verification was performed using simulated deterioration. The weight-factor matrix was trained on eleven cases of simulated deterioration, matching on-wing deterioration. The CEM was then verified on the remaining cases, both at take-off, climb and cruise conditions. The method was able to identify the simulation deterioration indicating that with this method it is possible to predict eight health indicators with six measurements. Especially the HIs of the HPC and HPT can be predicted with high accuracy, while the fan core + LPC efficiency and LPT flow capacity have a higher degree of uncertainty.

The CEM was validated using data from the Quick Access Recorder (QAR), which provides high-frequency (1 Hz) flight measurements. A data filtering procedure is applied to isolate near steady-state operating conditions, N1k settings close to the design-point, and secondary performance parameters, ensuring a homogeneous dataset suitable for reliable trend analysis. The method was evaluated on three separate cases. The first case consisted of a HPT failure, for which the CEM identified deterioration before other monitoring alerts were triggered. Furthermore, this case highlighted the added value of Companion Engine Analysis as residuals were used to find abnormal degradation compared to the companion engine. Companion Engine Analysis compares engines operating on the same aircraft to detect abnormal degradation, by assuming these engines operate under similar conditions. The second case consisted of proving the CEM was capable of not only predicting deterioration, but also predicting improved engine performance after a shop visit. In this case a HPT overhaul was performed, and this was shown in the rolling average of the HPT efficiency by a 2% increase, and the HPT flow capacity by a 1% increase. Finally, the compressor condition prediction accuracy was tested using waterwashes. Here, the CEM proved to accurately account for the effects of waterwashes, especially in the HPC, as the efficiency significantly improved after a waterwash. In addition, where the CEM requires about 2-5 seconds to calculate health indicators for a single data point, other methods such as genetic algorithms take hours, making the approach significantly more computationally efficient.

To conclude, the Component Exclusion Method combined with Companion Engine analysis has proved a novel approach which improves diagnostic accuracy of the GEnx-1B. CEM enables health estimation of underdetermined gas path cases (similar to the GEnx-1B), while avoiding high computational costs.

Preface

This master thesis is the result of my time as a student at the Faculty of Aerospace Engineering at TU Delft, and my year-long research at KLM Engine Services. I am sincerely grateful for the opportunity to carry out my thesis at KLM and hope to continue my professional career within the organization.

I would like to express my sincere appreciation to my supervisor, Wilfried Visser, whose guidance significantly deepened my interest in gas path analysis. His insightful discussions, interesting anecdotes during meetings, and support during challenging periods with GSPy were extremely helpful throughout this project.

I am also grateful to my supervisors at KLM, Tim Rootliep and Daniel Cisneros Acevedo, for their continuous guidance and support. Their mentorship not only contributed greatly to the quality of this work, but also made me feel welcome at KLM.

Furthermore, I would like to thank my colleagues at the KLM ES data team for their help and support. The walks to the coffee machine were very enjoyable, as were the conversations during lunch.

Finally, I would like to thank my friends and family for supporting me throughout my student time. Their encouragement has been essential to the completion of this work, and my Master Thesis.

Contents

List of Figures

List of Tables

Nomenclature

Acronym	Description
AM	Adaptive Modeling
ANN	Artificial Neural Network
ECBM	Engine Condition-Based Maintenance
CC	Combustion Chamber
CCS	Compartment Cooling system
CEOD	Continuous Engine Operating Data
CEM	Component Exclusion Method
CNN	Convolutional Neural Network
CEA	Companion Engine Analysis
DE	Differential Evolution
DL	Deep Learning
DP	Design-Point
EA	Evolutionary Algorithm
EEC	Electronic Engine Control
EGT	Exhaust Gas Temperature
EGTHDM	Exhaust Gas Temperature Hot Day Margin
FMU	Fuel Measuring Unit
GAN	Generative Adversarial Network
GPA	Gas Path Analysis
GE	General Electric
GSP	Gas Turbine Simulation Program
HDMR	High Dimensional Model Representation
HI	Health Indicator
HPC	High Pressure Compressor
HPT	High Pressure Turbine
HPTACC	High Pressure Turbine Active Clearance Control
IGV	Inlet Guide Vanes
KF	Kalman Filter
KLM	Koninklijke Luchtvaart Maatschappij
LPC	Low Pressure Compressor
LPT	Low Pressure Turbine
LPTACC	Low Pressure Turbine Active Clearance Control
LSTM	Long Short-Term Memory
ML	Machine Learning
MOPA	Multiple Operating Point Analysis
MRO	Maintenance, Repair and Overhaul
OD	Off-Design
OEM	Original Equipment Manufacturer
PHM	Prognostics and Health Management
RMSE	Root Mean Square Error
RNN	Recurrent Neural Network
RUL	Remaining Useful Life
TSFC	Thrust Specific Fuel Consumption
VBV	Variable Bleed Valve
VSV	Variable Stator Vane
WF	Fuel Flow
QAR	Quick Access Recorder

Greek Symbols

Symbol	Description
η	Efficiency
Δ	Change
ρ	Density
σ	Standard Deviation
μ	Mean
γ	Specific heat ratio

Roman Symbols

Symbol	Description
M	Mach number
W	Mass Flow
Y	measurement sensors
X	health indicators
P	Pressure
T	Temperature
N	Rotational Speed
r	residual

Subscripts

Symbol	Description
c	Case
c	Corrected
f	Fuel
h	Health Indicator
s	Static
s	Subset
t	Total
0,1,2,3,…	Engine Station Number

1

Introduction

In 2025, the International Air Transport Association (IATA) reported approximately one accident per 759,646 flights, with the majority of these incidents classified as non-critical events such as tail strikes, landing gear occurrences, or ground damage. Notably, only six fatal accidents were reported worldwide [1]. These statistics underscore the exceptionally high level of safety achieved in modern commercial aviation. Such performance is the result of decades of improvements in safety guidelines and technological advancements. Part of this safety record is the application of high maintenance standards, with Maintenance, Repair and Overhaul (MRO) organizations operating under strict regulatory guidelines and procedures. Included in this framework are aircraft engines, as engine reliability is essential to both safety and operational efficiency. Modern turbofan engines are complex systems operating under extreme heat and pressure, making them particularly susceptible to degradation. If undetected, this degradation may lead to increased fuel consumption, or in severe cases, to engine failures. Therefore, ensuring the health of an engine is a key part of aircraft operations.

Traditionally, airlines have relied on time-based or cycle-based maintenance schedules to ensure the reliability of their engines. While these methods are safe and well-established, they do not consider the actual health or degradation state of each engine. A more advanced approach is now emerging: On-condition maintenance scheduling. In this strategy, maintenance actions are based on the real-time condition of the engine rather than fixed intervals [2].

1.1. Context & Problem Statement

Within this approach, Engine Condition-Based Maintenance (ECBM) analyzes flight performance data to detect anomalies, diagnose degradation, predict remaining useful life (RUL) and propose required maintenance. A key method in this field is Gas Path Analysis (GPA), which estimates component health parameters based on measured engine performance. These parameters typically include component efficiency and flow capacity, which reflect physical degradation mechanisms such as fouling, erosion, and wear. This information can be used to adapt maintenance scheduling, adapt workscopes and detect anomalies [2].

At KLM Engine Services (ES), the shop has the ability to perform overhaul on four engine types, of which one is the GEnx-1B. This engine poses a particular problem for predictive maintenance approaches, due to its low number of gas path parameters. For accurate predictions, the more sensor measurements available the better. For the GEnx-1B, the OEM has limited the ability to install additional sensors [3]. Therefore, prediction of engine and component health has to be done using this limited sensor set, which presents a challenge. In addition, external factors such as changing operating conditions, measurement noise, and secondary performance parameters further reduce predictive accuracy.

Several methods have been proposed in previous studies to address these challenges, such as Multi-Operating Point Analysis (MOPA) and the integration of secondary performance parameters (SPPs) [3, 4, 5, 6]. In addition, companion engine analysis has been investigated at KLM, but not applied to health indicator calculation [7]. A recent advancement in the field of gas path analysis has enabled this research to improve these methods. A python-version of the Gas Turbine Simulation program (GSP)

has been published, allowing for personal adaptations and for computationally efficient health indicator simulations [8]. This enables the use of many more data points, effectively reducing influence of noise.

1.2. Research Objective & Question

The main objective of this research is to investigate a new GPA based method to predict health of components in the GEnx-1B engine. This method, called the Component Exclusion Method, utilizes a selection of health indicators called a subset to systematically run adaptive modeling simulations. Based on the research objective, the research question is as follows:

How can the Component Exclusion Method and Companion Engine Analysis be used to improve condition monitoring accuracy of the GEnx-1B with fewer sensors?

To answer this question, several subquestions are formulated:

1. *Which data sources (Snapshot, CEOD, QAR) are most relevant and reliable for implementing the Component Exclusion Method and Companion Engine Analysis?*
2. *How can GSPy be utilized and customized to accurately infer the health indicators of the GEnx-1B on-wing data?*
3. *How can the Component Exclusion Method be applied to improve current condition based maintenance approaches?*
4. *How can Companion Engine Analysis enhance diagnostic accuracy through cross-wing data comparison?*

1.3. Report Structure

The report is structured into three main parts. It begins with a literature study, followed by the methodology describing the Component Exclusion Method framework, and concludes with the verification and validation results. In chapter 2 the turbofan engine is introduced, with a focus on the GEnx-1B. This chapter also covers the different sources of engine data as well as deterioration present during on-wing operation. In chapter 3 the current state of predictive maintenance practices, both at KLM as in the industry, is explained. Then in chapter 4 modeling approaches are presented that are currently utilized in this field and could be used to improve the health indicator estimation of the GEnx-1B. Afterwards, chapter 5 describes the implementation platform used for this research. In chapter 6 the gas path analysis of the GEnx-1B is explained. Consequently, the Component Exclusion Method is described in chapter 7. A sensitivity analysis and the results of verification of the Component Exclusion Method are presented in chapter 8. Then, in chapter 9 the data processing is explained, and results of the CEM with on-wing data are presented. Finally, chapter 10 draws conclusions and chapter 11 lays out recommendations for future research.

Part I

Engine Health & Performance Literature Study

2

Turbofan Engines & GEnx-1B

This chapter introduces turbofan engine design, with a focus on the GEnx-1B engine. It begins with an overview of turbofan operation, including the Brayton cycle and the standard station numbering convention. Subsequently, engine deterioration is discussed, including key degradation mechanisms, their symptoms, and their impact on performance. An in-depth description of the GEnx-1B engine is then provided. Finally, the chapter presents an overview of the available engine data sources.

2.1. Turbofan Engine

The turbofan engine was developed to improve propulsive efficiency at high subsonic Mach numbers by reducing the exhaust jet velocity [9]. Unlike turbojet engines, a turbofan divides the incoming airflow into two separate streams. A portion of the airflow bypasses the engine core and is compressed solely by the fan before being accelerated through a separate nozzle. This flow path is referred to as the bypass stream or bypass region. The remaining airflow enters the engine core, where it passes through the Low-Pressure Compressor (LPC), followed by the High-Pressure Compressor (HPC), before entering the combustion chamber (CC). The high-energy combustion gases subsequently expand through the turbine stages, which provide the power required to drive the compressors and fan. An overview of a twin-spool turbofan engine is shown in Figure 2.1. In the sketch of the twin-spool turbofan, the two shafts are clearly visible. The first shaft, N1, connects the fan and LPC to the LPT. The second shaft, N2, connects the HPC to the HPT.

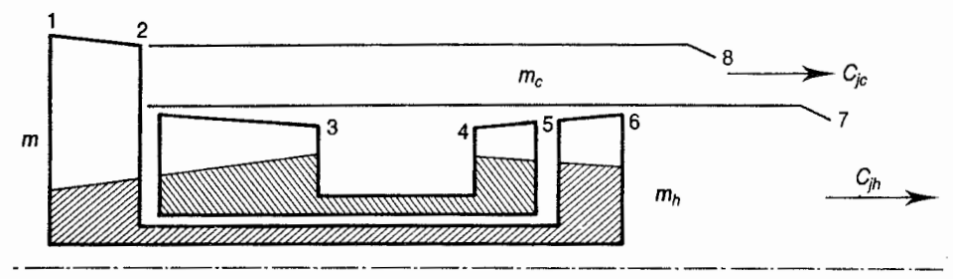

Figure 2.1: Twin-spool turbofan engine [9]

2.1.1. Brayton cycle

The thermodynamic processes occurring within a turbofan engine can be illustrated using the Brayton cycle, shown in Figure 2.2. In this diagram, temperature is plotted against entropy. The Brayton cycle consists of four idealized thermodynamic processes:

1. **Isentropic compression (0 $\rightarrow$ 3):** Ambient air enters the intake at station 0 and is compressed by the inlet and compressors. In Figure 2.2 the inlet included the fan. In addition, the compressor phase also consists of a low-pressure compressor (LPC) and a high-pressure compressor (HPC). This process increases the pressure and temperature while ideally maintaining constant entropy.

2. **Heat Addition (3 → 4):** In the combustor, fuel is injected and mixed with the compressed air. Afterwards, the mixture is burned at ideally constant pressure, causing a significant rise in temperature and enthalpy.
3. **Isentropic expansion (4 → 8):** The high-pressure gases expand through the high- and low-pressure turbines producing work to drive the fan and compressors. Ideally, this expansion takes place at constant entropy.
4. **Isobaric Heat Rejection (8 → 0):** After expanding through the turbines and nozzle, the exhaust gases leave the engine and transfer heat to the surroundings. In the ideal Brayton cycle, this process is represented as heat rejection at constant pressure, returning the working fluid to its initial ambient state. For aircraft gas turbines, this step is an open-cycle process since fresh ambient air continuously enters the engine while exhaust gases are discharged into the atmosphere.

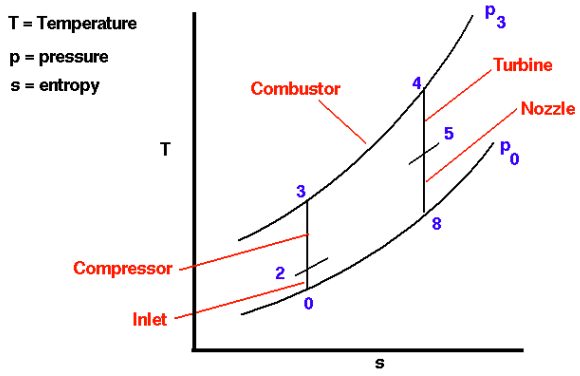

Figure 2.2: Ideal Brayton cycle for a turbofan engine [10]

The process depicted above is an idealized version of the Brayton cycle. In reality, pressure losses and inefficiencies lead to deviations from the ideal Brayton cycle, decreasing the total work done. However, this diagram does provide a clear framework, especially since the numbering is used in engine analysis.

2.1.2. Compressor maps

The simple 0D calculation of the station temperature and pressure can be done, but during off-design modeling the efficiency is not known. To address this, component maps can be used, like the axial compressor map in Figure 2.3. These maps provide a compact method of describing the off-design behavior of a compressor across a range of operating conditions.

A typical compressor map relates the pressure ratio to the normalized corrected mass flow. The x-axis represents the corrected mass flow, normalized with respect to inlet temperature and pressure. This allows the compressor map to be used for different operating conditions. The y-axis represents the pressure ratio across the compressor, one of the main performance indicators.

The curved lines correspond to the corrected shaft speeds (30%, 50%, 70%, 80%, 90%, 100%, 110%). Along these lines, the corrected shaft speed is constant, and shows how the pressure ratio changes with mass flow for a given speed. The grey dotted contours indicate regions of constant efficiency η_c (0.7, 0.75, 0.8, 0.82, 0.84, 0.86). The design point is also highlighted in the figure, and typically lies in the peak efficiency region. The design point is chosen such that it meets all performance requirements.

The operating line runs from the design point, at 100% corrected shaft speed, to idle conditions. This is how the engine operates during a full flight, and the operating region is carefully chosen not to enter unstable regions. The unstable regions consist of the surge/stall line, in the figure called the instability line, and the choke line. The choke line is not depicted in the figure but runs along the bottom of the compressor map, indicating when the pressure ratio is too low, while the corrected mass flow is high, the flow can reach sonic conditions, leading to choke.

In summary, compressor maps capture the relationships between pressure ratio, mass flow, spool speed, and efficiency. They provide a practical way to model component performance, especially at off-design conditions.

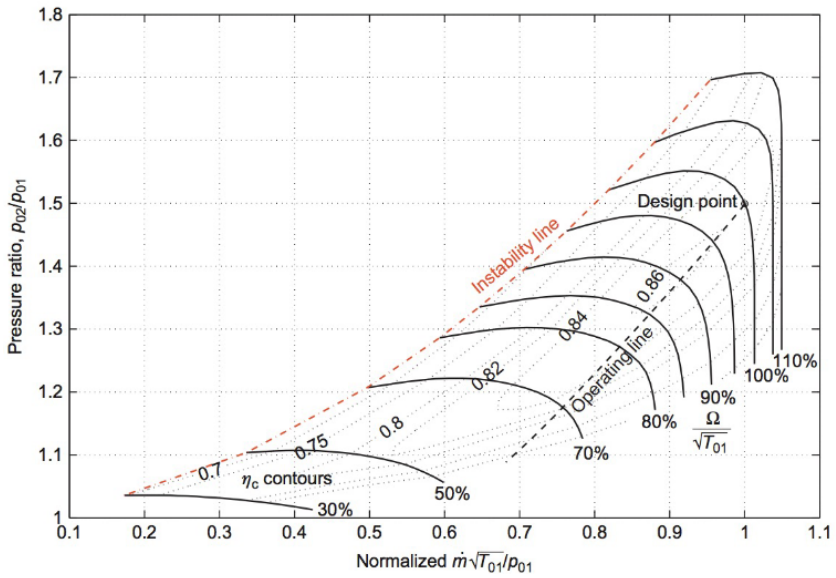

Figure 2.3: Axial compressor map [11]

2.1.3. Turbine maps

The same principle used for compressors can also be used for turbines. A turbine map, like the one in Figure 2.4, relates the turbine pressure ratio to the corrected mass flow. The curved lines represent the constant corrected shafts speed. Just like in compressor maps, the curves show how the turbine pressure ratio relates to the mass flow at constant shaft speeds.

The efficiency is also compared to the pressure ratio, typically shown on a second Y-axis on the right-hand side. The efficiency curves are shown in the top half of the chart. The lines represent constant corrected shaft speeds. At lower pressure ratios, the efficiency might not reach its optimum value, while at the design condition it does [11].

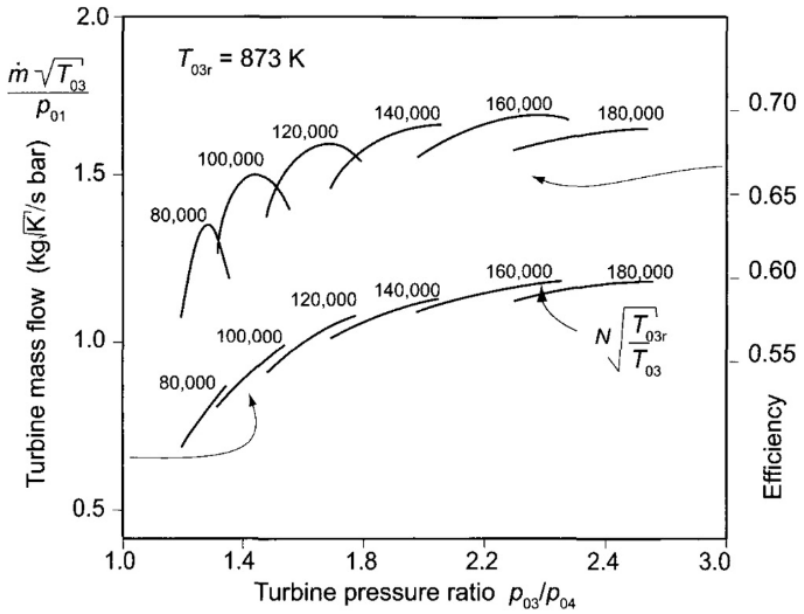

Figure 2.4: Turbine map [11]

In engine performance modeling, it is crucial that the components maps closely represent the actual engine. However, in many cases such maps are not available and must be estimated, scaled, or

derived from other maps. Given the strong nonlinearities in these maps, such approximations will introduce uncertainties, which will propagate throughout the performance model to eventually affect the accuracy of the predicted engine behavior.

2.2. Engine Deterioration

Throughout its operational lifetime, the gas path of an engine can be affected by a wide range of physical issues. Common examples include erosion, corrosion, fouling, and the accumulation of dirt. Engines may also suffer from foreign object damage (FOD), seal wear, and excessive tip clearances. In the turbine section, blades and stators can become burned, warped, or even partially missing. Additional problems may involve clogged fuel nozzles, as well as cracks in disks or blades caused by fatigue or operation beyond the intended limits [12]. These physical problems can affect the performance of the engine by changing the surface texture and the aerodynamic shape.

Erosion is defined as the removal of material from gas path components due to the impact of particles entrained in the airflow [13]. This deterioration mechanism typically increases the tip clearance in both compressors and turbines, resulting in a reduction in component efficiency. However, the effect on flow capacity differs between the two components. In compressors, the increased tip clearance leads to higher tip leakage losses, reducing the blades' ability to effectively compress the airflow and consequently decreasing the flow capacity. In turbines, the increased tip clearance reduces aerodynamic resistance, resulting in an increase in flow capacity [14].

Corrosion alters the blade surface texture and profile due to a chemical reaction between gas path components and contaminants in the inlet air or fuel [14]. This can lead to loss of material and an increase in surface roughness, resulting in a loss of performance. The effect is similar to erosion: A reduction in compressor flow capacity and efficiency, and a reduction in turbine efficiency but an increase in turbine flow capacity.

Fouling is defined as the accumulation of contaminants on gas path component surfaces [14]. Ambient air contains particles such as dust, dirt, sand, and other airborne contaminants that can adhere to blade and casing surfaces. This accumulation alters the surface roughness, blade geometry, and effective flow passage area of the components. Fouling has a significant impact on compressor performance, primarily causing a reduction in flow capacity and a decrease in efficiency. In turbines, fouling also reduces the effective flow area, leading to a decrease in flow capacity. Additionally, turbine efficiency is adversely affected due to the increased aerodynamic losses.

Foreign object damage (FOD) causes sudden and often severe deterioration of engine components, for example as a result of bird strikes or ingestion of debris. Depending on the nature and location of the damage, the flow capacity of both compressors and turbines may either increase or decrease. In all cases, however, component efficiency is adversely affected due to the disruption of the aerodynamic surfaces and flow field. If the damage is relatively minor, part of the performance loss may be recoverable through maintenance actions such as water washing [14].

Thermal distortion is in most cases located at the high pressure turbine where temperatures are highest [14]. It can lead to a change in aerodynamic shape of turbine blade, as well as crack growth and increased leakage. This results in an increased turbine flow capacity, and a decreased efficiency. In addition, if thermal distortion occurs in the compressor, both the flow capacity and efficiency of the compressor will decrease.

The effect of the aforementioned degradation causes can be found in Table 2.1, highlighting their physical effects and the resulting influence on the compressor and turbine. Each mechanism introduces different physical changes, but there are patterns in the compressor and turbine efficiency and flow capacity shifts. Although the degradation mechanisms can be investigated in great detail, this lies outside the scope of this study to provide an exhaustive analysis of their cause and progression. A clear understanding of these mechanisms is however important, as they form the basis for diagnosis and predictions of engine performance.

Table 2.1: Overview of degradation mechanisms, their effects, and performance impacts (derived from [5, 3, 14]).

Mechanism	**Degradation Effects**	**Compressor W_c shift**	**Compressor η shift**	**Turbine W_c shift**	**Turbine η shift**
Erosion	Enlarged tip clearance; Rougher texture; Reduced surge margin	↓	↓	↑	↓
Corrosion	Surface deterioration; Blade profile	↓	↓	↑	↓
Fouling	Surface deposits increasing roughness; Blade geometry alterations; Flow passage obstruction; Lower surge margin	↓	↓	↓	↓
FOD	Immediate performance disruption; Possible blockage or inoperability	↓ / ↑	↓	↓ / ↑	↓
Thermal Distortion	Changes in surface area; Increased leakage	↓	↓	↑	↓

2.3. GEnx-1B Engine

This section will introduce the layout of the GEnx-1B engine, which is central to this research. The airflow and engine control systems will be discussed, as well as performance enhancing system such as the active clearance control systems.

2.3.1. Introduction

In 2014, Air-France KLM selected the GEnx-1B to power its Boeing 787 fleet [15]. This agreement also included a maintenance, repair and overhaul (MRO) partnership, enabling Air France-KLM to provide MRO services for the GEnx-1B.

The GEnx-1B is a high bypass twin-spool turbofan developed by General Electric and is part of the next generation of GE engines focused on fuel efficiency, noise reduction and lower emissions. This is done by incorporating a high bypass ratio of 9.0, a high pressure ratio of 43.8 and a massive fan with a diameter of 2.8 meters [15]. In addition, it incorporates several control systems to ensure optimal performance. These play a central role in regulating airflow, cooling and stability.

2.3.2. Primary Airflow Control Systems

The primary airflow control subsystem of the GEnx-1B manages airflow within the compressor section. This subsystem includes the Variable Stator Vane (VSV) actuators, the Variable Bleed Valve (VBV) actuators, both of which are hydraulically controlled by the Fuel Metering Unit (FMU). The FMU varies hydraulic pressures in response to commands from the EEC in order to position the actuators [16].

The Variable Stator Vane (VSV) system consists of the compressor inlet guide vanes (IGVs) and variable stators vanes of stages 1 through 4 of the compressor. These vanes are connected to circumferential actuation rings through level arms, allowing them to rotate about their radial axis. The VSV system adjusts the airflow angle entering the compressor stages, thereby optimizing airflow. This is vital for maintaining surge margin at different power settings and ambient conditions.

The Variable Bleed Valve (VBV) system consists of ten doors located aft of the booster, which interconnect the compressor inlet flow path with the fan discharge stream. The FMU adjusts hydraulic pressure to move the actuators and position the bleed doors according to the schedule commanded by the EEC. The VBVs are primarily used to manage airflow during engine start, acceleration, and transient conditions to prevent compressor surge and maintain stable operating.

2.3.3. Performance Enhancement Systems

Besides airflow control systems, the GEnx-1B also employs several performance enhancement systems to increase the efficiency and operability of the engine.

The High-Pressure Turbine Active Clearance Control (HPTACC) is one of the systems which is added to increase efficiency. The purpose of this system is to maintain optimal HPT blade tip clearance. This is achieved by regulating cooling airflow onto the exterior of the HPT case. As the turbine case cools, it contracts slightly, reducing the clearance between the shrouds and rotor blade tips. Minimizing this gap reduces tip leakage losses and increases turbine efficiency, directly improving the engine's specific fuel consumption (SFC). The HPTACC valve is controlled by the EEC and an electro-hydraulic servo valve [16].

The Low-Pressure Turbine Active Clearance Control (LPTACC) system performs a similar function to the LPT. It controls the clearance between the LPT rotor blades and surrounding shrouds by adjusting cooling airflow to the turbine case. Like the HPTACC, tighter clearances enhance aerodynamic efficiency and reduce SFC, particularly during cruise when the turbine operates at steady-state temperatures.

The Transient Bleed System (TBS) is designed to improve compressor stability during starting and engine acceleration from idle, and in icing conditions. The system achieves by unloading the compressor, which increases the stall margin. The Transient Bleed Valve (TBV) is located on the aft compressor case and opens to discharge compressor air into the bypass stream. By reducing pressure in the compressor core during acceleration, the TB system ensures reliable operation.

Finally, the core compartment cooling system (CCS) can provide cooling airflow to the engine's under cowl environment. The system uses fan discharge and VBV air, which passes through a valve controlled by the EEC. The valve regulates airflow depending on engine power setting and ambient conditions.

2.4. Engine Data

This section presents the available data sources for the GEnx-1B engine. Three distinct data sources are considered and discussed in detail, followed by a comparative summary highlighting their respective advantages and limitations. Additionally, the accuracy of the associated measurement instruments is addressed.

2.4.1. Data Sources

Modern aircraft are equipped with multiple systems to record engine and flight parameters. These data sources include snapshot measurements, data from the Engine Monitoring Unit (EMU), also called Continuous Engine Operating Data (CEOD), and data from the Quick Access Recorder (QAR). Snapshot data gives a single measurement of ambient conditions and gas path parameters three times during a flight, in take-off, climb and cruise. The parameters are recorded in a short time interval (3 to 30 seconds), aggregated and stored or sent back to the ground in real-time using ACARS (Aircraft Communication Addressing and Reporting System). To visualize the snapshot data output, the output matrix is given below.

$$Y = \begin{bmatrix} N_1^{(\text{TO})} & M^{(\text{TO})} & T_{t2}^{(\text{TO})} & P_{t2}^{(\text{TO})} & T_{t25}^{(\text{TO})} & T_{t3}^{(\text{TO})} & P_{s3}^{(\text{TO})} & \dot{m}_f^{(\text{TO})} & N_2^{(\text{TO})} & T_{t49}^{(\text{TO})} \\ N_1^{(\text{Climb})} & M^{(\text{Climb})} & T_{t2}^{(\text{Climb})} & P_{t2}^{(\text{Climb})} & T_{t25}^{(\text{Climb})} & T_{t3}^{(\text{Climb})} & P_{s3}^{(\text{Climb})} & \dot{m}_f^{(\text{Climb})} & N_2^{(\text{Climb})} & T_{t49}^{(\text{Climb})} \\ N_1^{(\text{Cruise})} & M^{(\text{Cruise})} & T_{t2}^{(\text{Cruise})} & P_{t2}^{(\text{Cruise})} & T_{t25}^{(\text{Cruise})} & T_{t3}^{(\text{Cruise})} & P_{s3}^{(\text{Cruise})} & \dot{m}_f^{(\text{Cruise})} & N_2^{(\text{Cruise})} & T_{t49}^{(\text{Cruise})} \end{bmatrix}$$

Secondly, CEOD is a stream of data consisting of a 697 engine parameters recorded during the flight. The frequency of CEOD lies around 10 Hz. The first recorded GEnx-1B flight with CEOD data was around mid-2016. This makes it very useful for vibration analysis for example, where high frequency data is preferred. However, flight recorders can dynamically adapt their sampling frequency, and therefore the time series of measurements can have very different frequencies, and they can vary throughout the flight.

Another type of data stems from the Quick Access Recorder (QAR) which records 5847 aircraft flight parameters at a frequency of 1Hz [17]. The data is easily accessible using memory cards, USB or

cellular network. However, this data only contains flights since 2018. A visualization of data from the EMU (CEOD) and data from the QAR is given below.

$$Y = \begin{bmatrix} N_{1,t} & M_t & T_{t2,t} & P_{t2,t} & T_{t25,t} & T_{t3,t} & P_{s3,t} & \dot{m}_{f,t} & N_{2,t} & T_{t49,t} \\ N_{1,t+1} & M_{t+1} & T_{t2,t+1} & P_{t2,t+1} & T_{t25,t+1} & T_{t3,t+1} & P_{s3,t+1} & \dot{m}_{f,t+1} & N_{2,t+1} & T_{t49,t+1} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ N_{1,t+n} & M_{t+n} & T_{t2,t+n} & P_{t2,t+n} & T_{t25,t+n} & T_{t3,t+n} & P_{s3,t+n} & \dot{m}_{f,t+n} & N_{2,t+n} & T_{t49,t+n} \end{bmatrix}$$

A summary of the different types of engine data is given below.

Table 2.2: Comparison of engine data

Model	Availability	Advantages	Limitations	Use Cases
Snapshot	98% of flights - since 2015	Data available since first flight; Can capture long term effects	Only 3 data points per flight	Long-term analysis
EMU (CEOD)	30% of flights - since 2016	High frequency data; Large dataset	Missing data; Variable sampling frequency	High frequency analysis (vibration)
QAR	80% of flights - since 2018	Comprehensive coverage; Easy accessible; Large dataset	No data before 2018	Short & Long-term analysis

2.4.2. GEnx-1B Sensor Availability

One of the primary challenges in Gas Path Analysis (GPA) modeling of the GEnx-1B engine is the limited number of available sensors compared to earlier engine types. For instance, the CF6-80E1 engine is equipped with ten gas path measurements, whereas the GEnx-1B is equipped with only six. This reduced sensor set restricts the number of measurable parameters, thereby increasing the level of uncertainty in performance and health estimation.

A schematic of the GEnx-1B engine with its measured parameters is shown in Figure 2.5. The six available gas path measurements for the GEnx-1B engine are:

- $T_{t,25}$: The HPC inlet total temperature
- $T_{t,3}$: The HPC outlet total temperature
- $P_{s,3}$: The static HPC outlet pressure
- W_f: The fuel flow
- N_2: The core shaft speed
- $T_{t,49}$: The Exhaust Gas Temperature (EGT)

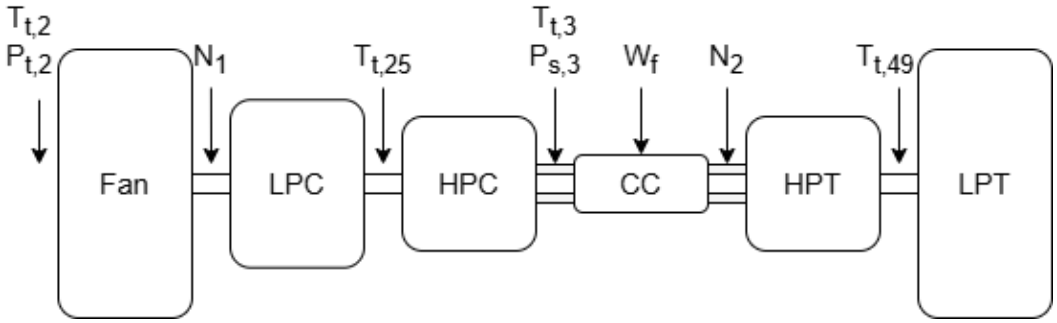

Figure 2.5: GEnx-1B sensor placement

In addition to these gas path measurements, the GEnx-1B engine is equipped with two vibration sensors that provide critical information for detecting mechanical faults. The first is the number 1 bearing

accelerometer, which is mounted on the number 1 bearing housing flange. It measures the vibration amplitude and frequency of the fan section and transmits this information to the Engine Monitoring Unit (EMU) [16]. The unit of measurement is Mils S.A., where one Mil equals one-thousandth of an inch and S.A. stands for Single Amplitude—a measure of the peak displacement of vibration.

The second vibration sensor is located in the Turbine Center Frame (TCF). This accelerometer measures the vibration characteristics of the engine core and is also connected to the EMU. Its output is expressed in IPS pk (Inches Per Second, peak value), representing the maximum vibration velocity of the turbine structure. These vibration measurements are vital for identifying imbalance, bearing wear, and other mechanical issues that may not be directly detectable through gas path parameters alone [18].

Like all measurement systems, the sensors equipped on the GEnx-1B are subject to noise, bias, and calibration drift. Quantifying the accuracy and precision of each instrument is therefore essential to ensure reliable diagnostic results. Table 2.3 summarizes the key characteristics and measurement accuracies of the GEnx-1B sensors. This information forms the foundation for uncertainty quantification in GPA modeling and subsequent health estimation analyses.

Table 2.3: Accuracy of GEnx-1B sensors [16, 3]

Parameter	**Sensor type**	**Instrument Range**	**Recommended system accuracy**	**Unit**
Shaft Speed - N1, N2	Magnetic Induction	0 - 14,790	$\pm$ 1%	rpm
Total Temperature - $T_{t,2}$, $T_{t,25}$, $T_{t,3}$, $T_{t,49}$	Thermocouple	-75 - 1200	$\pm$ 10$^{\circ}C$	$^{\circ}C$
Fuel Flow - W_f	Mass fuel flow meter	400 - 30,000	$\pm$ 3.5%	pph
Total Pressure - $P_{t,2}$	Pressure transducer	1.6 - 20	$\pm$ 0.01	psia
Static Pressure - $P_{S,3}$	Pressure transducer	5 - 750	$\pm$ 0.36	psia
Vibration Fan	Accelerometer	40	$\pm$ 0.1 Mil	Mils S.A.
Vibration Core - TCF	Accelerometer	20	$\pm$0.1 IPS	IPS pk.

3

Maintenance Concepts

Predictive maintenance has emerged as a key strategy in modern maintenance operations. In contrast to traditional time- or cycle-based approaches, it utilizes operational data and degradation models to optimize maintenance planning, thereby reducing downtime and associated costs. This chapter introduces key maintenance concepts, with a focus on Engine Condition-Based Maintenance (ECBM). The principles of diagnostics and prognostics are discussed, followed by an overview of current practices at KLM Engineering & Maintenance (E&M) and relevant developments within the industry.

3.1. Engine Condition-Based Maintenance

Engine Condition-Based Maintenance (ECBM) is a maintenance strategy, where the choice of maintenance depends on actual component/engine health [19]. While the standard maintenance strategy is mainly focused on time-based or cycle-based maintenance, ECBM can be viewed as a means to mitigate failure risk while minimizing engine down-time [20].

3.1.1. Methodology

ECBM consists of three steps: data acquisition, data processing and maintenance decision-making. It attempts to avoid unnecessary maintenance tasks by selecting maintenance actions only if there is a high probability of a component fault [21].

Data acquisition consists of the collection and storage of measured data. This usually includes gas path parameters, as well as vibration data. The next step is data processing, which focuses on selecting relevant parameters and filtering those parameters for noise, for instance using a Kalman filter. Then Health Indicators (HI) are created using AI techniques or GPA models. These HI's give a representation of the health of the engine, or of a certain component. Finally, according to degradation trends of the HI's, the RUL is predicted which can then influence maintenance schedules.

From a maintenance and fleet perspective there is a split in methodology between two engine health management terms: Diagnostics and Prognostics. Diagnostics entails the immediate fault detection and identification [20]. From this perspective, reductions in operating costs can be achieved by avoiding unscheduled events, as well as identifying malfunction of non-critical subsystems, such as fuel leaks, which can lead to a lower efficiency. Prognostics, although related, deals with predicting the state of the engine and its components. This is usually portrayed in the form of degradation trends. Included in these trends are wear related processes, explained in section 2.2. The prognostics models use engine data as well as maintenance events to provide an expected Remaining-Useful-Life, RUL. In Figure 3.1 an overview is given of the current MRO practices.

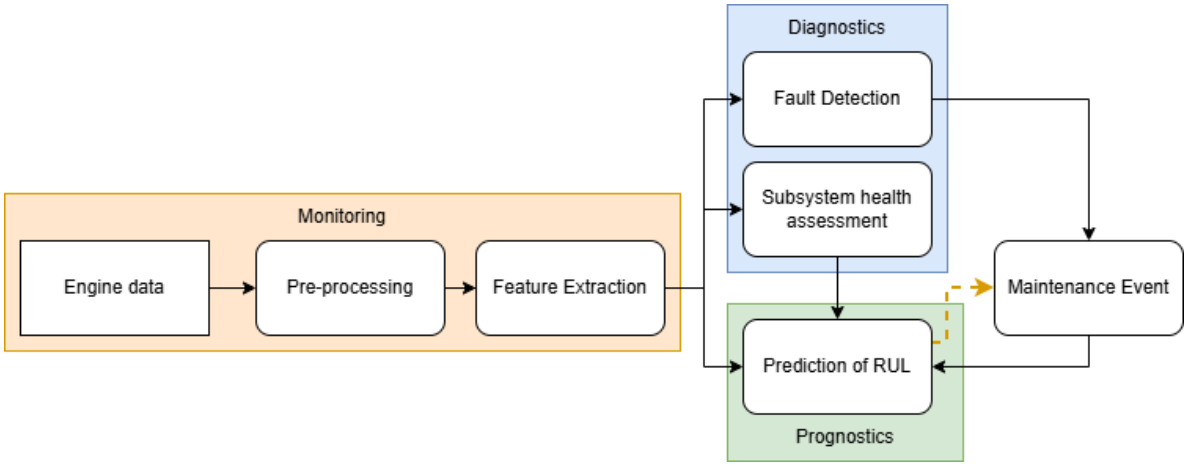

Figure 3.1: Overview of engine diagnostics and prognostics (adapted from ([22], [23]))

3.1.2. KLM Maintenance Practices

KLM currently employs a system known as PROGNOS for initial prognostic monitoring [5]. PROGNOS utilizes engine data to track trends in measured parameters over time. A rule-based warning mechanisms is then applied to signal when specific thresholds are exceeded, indicating potential component failures. Although PROGNOS has proven to be valuable, such as the identification of a HPT failure, the system is challenging to tune, as the rules generate many false positives. These false alarms undermine the systems reliability, and therefore it is not fully adopted as a prognostics model.

MRO companies can benefit from ECBM due to the Power by the Hour (PBH) contracts they give out to operators, in addition to decreasing the Total Cost of Ownership (TCO) of their own fleet. The PBH is a contract where a client pays a per-flight-hour rate for maintenance, repair and overhaul. This model is profitable for both parties, since it offers a more constant stream of income, while the client has monthly expenses, rather than a large sum at once. By using ECBM, the MRO facility can increase the flight hours of the engines they maintain, increasing their revenue [6].

Although ECBM is currently not directly applied to customer engines, several practices can be implemented by KLM Engine Services to lower the Total Cost of Ownership (TCO). First, condition monitoring can be used to extend bench time. Bench time is the desired number of cycles between maintenance events. For example, the estimated bench time for the GEnx-1B is 3000 cycles, but increasing the accuracy of this prediction could enable more reliable shopvisit scheduling and reduce the frequency of unscheduled removals.

In addition, during a shop visit KLM can perform additional maintenance activities belong the required scope if doing so leads to a reduction in long-term cost per flight hour [6]. For example, if a maintenance activity for the HPT is planned but condition-monitoring shows that the LPT has also degraded, an additional LPT maintenance activity could be performed to ensure long term health of the LPT.

Furthermore, in some cases it could be more economical to remove the engine earlier if this prevents major component scrap or catastrophic failure. If the HPT has a low engine health for example, failure of a blade will also damage the LPT, even if the LPT module has a high health indicator value. Removal and maintenance before the planned maintenance event could prevent this failure.

Finally, by gaining insight into component health prior to engine disassembly, KLM ES can better plan the shop visit work scope, order parts in advance and assign labor more efficiently. This predictive planning shortens turnaround time and reduces unnecessary inspection. All these consequences lead to a cost reduction. The overall goal of the current ECBM planning for the next two years can be seen in Figure 3.2, where component-wise engine health predictions are used to support maintenance workscoping decisions.

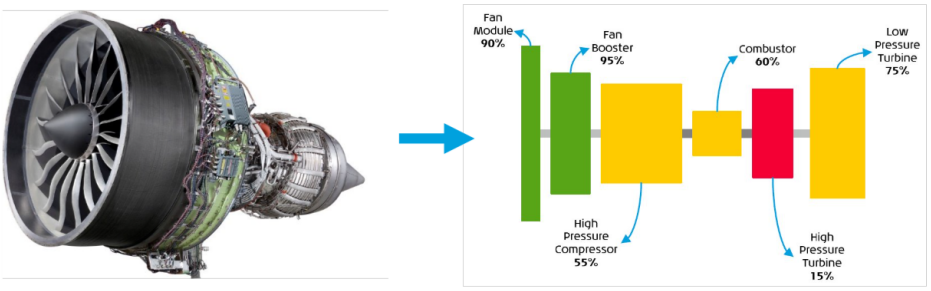

Figure 3.2: KLM future ECBM goal

A vast amount of research has been done on this topic by students within KLM, shown in Figure 3.3. This was primarily focused on the modeling of turbofan engines. Starting with the CF6 and CFM, research was later expanded to the newer GEnx-1B engine. Moorselaar [24] started with the GSP modeling of the GEnx engine. After that, Rootliep [3] and Otten [4] continued his work, improving the GEnx model. de Bruin [5] focused on reducing computational time using a High-Dimensional Model Representation (HDMR). Then Brachmi [6] extended his work by implementing secondary performance parameters. More recently, Fokkema [25] and Cisneros Acevedo [26] have taken a more data driven approach. This was made possible due to the vast amount of data currently available. Finally, Singhvi [7] conducted an analysis on companion engines. The work in this thesis extends on the work by Singhvi, using companion engine analysis to improve predictive maintenance of the GEnx-1B.

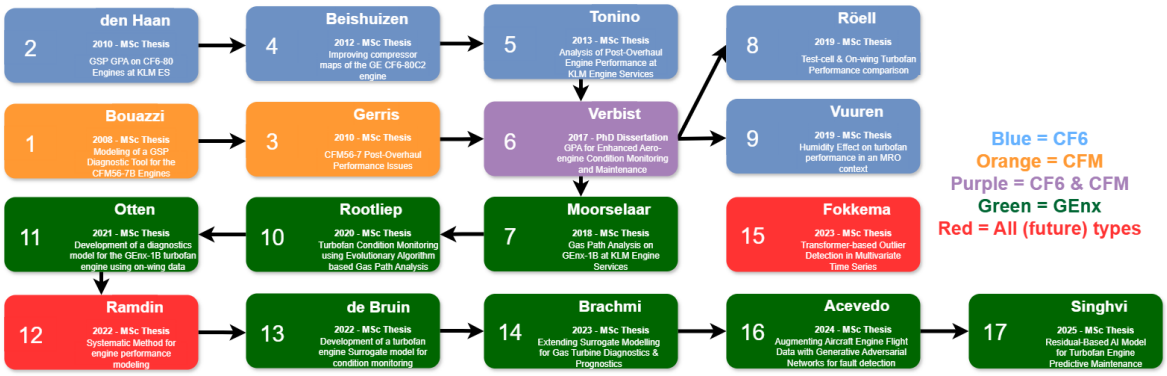

Figure 3.3: Previous master thesis projects at KLM Engine Services on engine condition based maintenance

3.2. Industry Trends

Diagnostic practices on aircraft engines are as old as the engines themselves. First monitoring abnormal levels in exhaust gas temperature (EGT), spool rotational speed and engine vibration level, this was the beginning of engine health monitoring. Since reliability and safety became an important constant, engine gas path parameters were added to give the operator more information about the state of the engine. At the end of the 20th century, electronic hardware was installed on all wide body aircraft [20]. This first system, airborne integrated monitoring system (AIMS) consisted of engine control parameters, and later also gas path parameters. The goal of AIMS was to record exceedance from stored thresholds during flight. Over time, the role of diagnostics and prognostics has changed from maintenance troubleshooting to a more proactive role. Since the introduction of PBH contract, as explained before, the inclusion of prognosis in maintenance strategies gives a win-win situation for both the EOM, and the operator [20].

The last decade, due to the rise of Artificial Intelligence, there has been a shift to data-driven predictive maintenance. Integrating fleet-wide data has been used to benchmark individual engines. This allows the identification of degrading engines by comparing to the fleet benchmark [27]. For example, Rolls Royce uses a system named Engine Health Management to track the health state of their engines

[28]. The development of predictive maintenance has also led to the replacement of static exceedance limits with adaptive alarm thresholds, which reduce false positives under variable operating conditions [28]. With the introduction of high-frequency engine condition monitoring systems such as QAR and EMU data and increased availability of fleet-wide operational data, more sophisticated anomaly detection and fault isolation techniques have been implemented. The focus is increasingly shifting toward combining physical knowledge with data-driven insights, enabling more robust decision-making under uncertainty. Overall, the industry trend has moved from threshold-based monitoring toward predictive and prescriptive maintenance, with a strong emphasis on integrating data-driven methods, adaptive modeling, and cross-fleet analysis. This evolution provides the foundation for the research presented in this thesis, which aims to leverage both physics-based and data-driven approaches to improve the accuracy and reliability of gas turbine engine health assessments.

4

Diagnostics & Prognostics

Several modeling approaches can be applied within diagnostic and prognostic frameworks for predictive maintenance of turbofan engines. This chapter presents a literature review of physics-based approaches and their associated limitations, followed by an overview of data-driven techniques and companion analysis methods. Finally, a trade-off analysis is performed to identify the most suitable methodology for the present research.

The primary objective of physics-based approaches is to estimate health indicators from available engine measurements. Since these health indicators cannot be measured directly, they must be inferred through reverse engineering of the engine gas path. Such approaches typically rely on tools such as the Gas Turbine Simulation Program (GSP), often combined with state-estimation techniques including the Kalman Filter. In contrast, data-driven approaches employ machine learning and artificial intelligence methods to identify patterns and relationships within operational data.

4.1. Physics-based Approaches

The most common approaches are the physics-based approaches. These methods rely on physical relationships like thermodynamic laws, conservational laws and component performance models to find the gas path parameters in an engine. By linking health indicators to sensor measurements, they can highlight engine component degradation. In addition, the strength of these approaches lies in their transparency, as the output of these models can be trailed back to an input. However, they often require detailed engine knowledge, tuning, and can struggle with non-linear behavior. In this section, the main physics-based approaches are explained.

4.1.1. Exhaust Gas Temperature (EGT)

A simple way to model the state of the engine is by looking at the EGT. The idea behind this is: The hotter the engine has to run for the same performance, the more deteriorated it is. Although this is true in most cases, using only the EGT does not allow one to locate the degradation and quantify it. The EGT Hot Day Margin (EGTHDM) is commonly used in industry as a health indicator and represents the difference between the maximum allowable EGT and the measured EGT during operation, corrected for hot-day operating conditions. As engine deterioration increases, the available margin decreases because a higher turbine inlet temperature is required to produce the same thrust. A low EGTHDM is therefore an indication that engine performance has degraded and that maintenance actions may be required in the future.

4.1.2. Classical GPA

The classical GPA method relies on a Fault Coefficient Matrix (FCM) method, which relies on measuring changes in dependent variables and comparing them with tables of calculated expected deviations [12]. It works for single faults but struggles with multiple simultaneous faults because different fault combinations can produce similar parameter shifts. A typical influence matrix can be derived to define a set of differential equations which interrelate engine performance parameters. It is applicable to

multiple simultaneous faults and matrices can be inverted to directly solve for independent parameter changes from measurements. It is done by rearranging into a set of linear equations. This means it is also computationally simple, and applicable to different engine types. However, it does require custom baselines for every engine, and separate coefficient sets have to be used for take-off, climb and cruise.

4.1.3. Gas Turbine Simulation Program (GSP)

The first models of the gas turbine consisted of only manual cycle and component performance and structural stress calculations [29]. Later, detailed modeling methods such as Computational Fluid Modeling (CFM) were used to calculate engine parameters. This was combined with the *zooming* concept, which refers to a multiscale modeling approach, where the entire engine is simulated using simplified physics but small section are modeled using complex physical models.

The engine simulation programs available are the Gas turbine Simulation Program (GSP) [29], Gasturb [30], NPSS [31] and PROOSIS [32]. Due to the long affiliation of KLM and TU Delft with GSP, this program will be used as simulation tool out of the above mentioned. GSP is a generic, object-oriented simulation environment designed for comprehensive gas turbine performance analysis. It uses the 0-D calculation framework, where averaged values are taken at point in the gas path for the temperature and pressure. In addition, GSP simulates the gas turbine cycle by solving sets of non-linear differential equations. These equations are based on the fundamental conservation laws (mass, energy, momentum). GSP then assumes a steady-state process to determine equilibrium operating points. To quantify the relationship between mass flow, shaft speed, efficiency and flow capacity, component maps are used. They are represented as look-up tables, and the accuracy of these tables is vital to ensure good engine representation.

The solution procedure is based on the multi-variable Newton-Raphson numerical method [29]. The objective is to find the values of state variables S that all error variables E become zero. This starts by an initial guess of the state variables, which is often the design point. Then the error is calculated and a first order Taylor approximation is used to linearize the error around the current guess:

$$\Delta E = J \cdot \Delta S \tag{4.1}$$

Where J is the Jacobian matrix. Using each state variable, the Jacobian matrix is built. This is done by changing each state variable slightly and computes the change in error, which leads to a $\delta e_j / \delta s_j$. The next step consists of inverting the matrix, and solving for the correction, seen in Equation 4.2.

$$\Delta S = J^{-1} \cdot \Delta E \tag{4.2}$$

The state variables can be updated using a step size factor to control stability:

$$S_{i+1} = S_i - f \cdot J^{-1} \cdot E_i \tag{4.3}$$

Finally, a convergence check is done to see if all values of E_i are below a certain tolerance.

For transient operation, GSP employs a quasi-steady-state simulation approach using the modified Euler method. In addition, GSP is particularly well suited for off-design analysis, in which the operating point of a specific engine configuration is determined under varying conditions, including ambient conditions, power settings, and engine health states. However, this capability is strongly dependent on the availability and accuracy of component performance maps. Original Equipment Manufacturer (OEM) component maps are generally proprietary and therefore unavailable for public use. As a result, component maps must typically be estimated by scaling generic maps to a specified design point [27]. Recent developments, such as the application of differential evolution optimization techniques for map tuning and polynomial scaling methods, have demonstrated promising improvements in off-design prediction accuracy [33].

4.1.4. GSPy

GSPy is a Python-based adaptation of the Gas Turbine Simulation Program (GSP), developed using the same procedures as GSP, but in an open-source and more flexible environment, Python [8]. This enables easy integration and customization for different research applications.

4.1.5. Adaptive Modeling

Another Gas Path Analysis (GPA) method implemented in GSP is adaptive modeling. This approach enables the estimation of component deterioration and fault conditions from available sensor measurements by adapting component characteristics within the engine model. Two principal approaches exist for constructing an adaptive gas turbine model. The first employs an external numerical iteration loop around the existing engine model. The second extends the internal equation set of the gas turbine model by incorporating additional equations that enforce agreement with measurement data through adaptation of the unknown variables representing component condition deviations. GSP adopts the latter approach for its adaptive modeling methodology [29].

4.1.6. Multi-Operating Point Analysis (MOPA)

To address the underdetermined system, Visser and Stamatis et al. [29, 34] suggest the use of multi-point measurements. By incorporating data from various operating points, additional constraints can be introduced, effectively turning an underdetermined system into a solvable one. This concept forms the foundation of Multi-Operating Point Analysis (MOPA), which uses the assumption that degradation remains constant across different operating points to provide a more robust estimation of health indicators. The accuracy of these models also relies on the tuning of the baseline model, often using calibration factors to compensate for model errors [29]. Rootliep et al. [3] also states that multi-point GPA improved robustness over single-point GPA.

4.1.7. GPA with Health Indicator Selection

A method for mitigating the underdetermined problem is through the selection of health indicators [35][36]. Both papers propose methods that, while differing in some details, share a similar framework: First, all possible combinations of health indicators are defined based on the amount of measurement parameters. Adaptive modeling is then performed on each health indicator combination independently. The results from each combination are processed and combined to create a fault signature which indicates the engine component most likely to be deteriorated.

This approach has been shown to be effective in accurately identifying faulty components, provided that the selected subsets of health indicators include all relevant root causes [36]. However, determining which subsets to use is typically performed manually, which can be a time-consuming process, especially when dealing with the large datasets considered in this research. This will have to be automated to be applicable of this research.

4.1.8. Residual and Filtering Methods

Residual methods are used to calculate degradation in a different way. Instead of calculating the health indicators, GSP is used to form a baseline of gas path parameter data. The real-time engine data is then compared to the parameters of this baseline and the residuals are used as degradation indicators [27]. This is also used in sister/companion engine analysis by Singhvi [7]. A limitation of residual methods is that the accuracy depends strongly on the fidelity of the baseline model.

In order to track component degradation over time, filtering techniques can be employed to reduce the influence of measurement noise [27, 37]. One commonly applied method is the Kalman Filter (KF), which estimates the state of a variable over time and predicts its future evolution. Kalman Filters are particularly well suited for gas turbine prognostics due to their ability to effectively handle noisy measurements while maintaining relatively low computational complexity [5]. However, their application also presents several challenges. Accurate prognostic performance may require the implementation of more advanced and computationally demanding KF variants. In addition, Kalman Filters are sensitive to smearing effects, where degradation associated with one component may be incorrectly attributed to another component within the engine model.

4.2. Limitations

A fundamental challenge in deterministic Gas Path Analysis (GPA) for the GEnx-1B engine is the underdetermined nature of the diagnostic problem. Specifically, the number of available sensor measurements, represented by Y, is smaller than the number of unknown health indicators, represented by X, resulting in the absence of a unique solution. The measured gas path parameters of the GEnx-1B can

be seen in Equation 4.4. The available sensor set consists of six measurements: The total temperature at the fan outlet $T_{t,25}$, the total temperature at the HPC inlet $T_{t,3}$, the static pressure at the HPC inlet $P_{s,3}$, the fuel mass flow $\dot{m}_f$, the core speed N_2, and the total temperature at the HPT inlet $T_{t,49}$. Consequently, the measurement space has a dimension of six:

$$Y = \left[\, T_{t,25},\ T_{t,3},\ P_{S,3},\ \dot{m}_f,\ N_2,\ T_{t,49} \,\right], \quad \dim(Y) = 6 \tag{4.4}$$

In contrast, the health indicator vector contains eight unknown component condition parameters:

$$X = \left[\Delta\eta_{\text{fan,LPC}},\ \Delta W_{c_{\text{fan,LPC}}},\ \Delta\eta_{\text{HPC}},\ \Delta W_{c_{\text{HPC}}},\ \Delta\eta_{\text{HPT}},\ \Delta W_{c_{\text{HPT}}},\ \Delta\eta_{\text{LPT}},\ \Delta W_{c_{\text{LPT}}}\right], \quad \dim(X) = 8 \tag{4.5}$$

Since $\dim(Y) < \dim(X)$, the system of equations linking measurements to health indicators is underdetermined [33]. This states that there are multiple solutions possible for solving the equation set. This represents a fundamental limitation of the problem and must be considered when selecting an appropriate diagnostic approach.

4.3. Data-driven Approaches

Data-driven approaches have emerged as a powerful tool in engine health diagnosis and prognosis. Rather than relying on physical relationships between components as used in physics-based approaches, data-driven approaches leverage large datasets of sensor data to identify patterns. Using these patterns a prediction can be made about component degradation. Furthermore, the increasing availability of engine sensor and maintenance data combined with advances in machine learning have enabled great advances in accuracy of these methods. While data-driven methods often face challenges like noise, model inaccuracy and data scarcity, they provide significant opportunities to increase current diagnostics and prognostics. This section will portrait the most promising of these approaches.

4.3.1. Surrogate Models

Although GSP has proven to be very effective in simulating turbofans, its computational cost becomes a disadvantages when dealing with large-scale tasks, such as large datasets, iterations or optimizations. A common solution is the use of surrogate models, which approximate the output of GSP with an increase in computational speed and minor loss of accuracy. These surrogate models typically consist of machine learning architectures that can be trained on representative datasets generated using GSP, test-cell measurements, or on-wing measurements [26].

Within KLM Engine Services (ES), a surrogate model for the GEnx-1B has been developed using data produced by a high-fidelity GSP digital twin. Approximately 1.2 million simulated scenarios were generated, covering a broad range of operating conditions [26]. The resulting dataset provided corresponding gas path parameters. These outputs were mapped directly from the GEnx-1B turbomachinery maps, which had been calibrated against single engine test-cell correlation data. Consequently, the surrogate model inherits a bias toward the reference test-cell engine and may not perfectly generalize across an entire fleet [26]. The model consists of a four-layer fully connected neural network (FCNN). The internal architecture begins with an initial layer of 64 units, followed by layers that progressively halve until the last layer, which is the same size as the number of sensors [26]. A Min-Max scaler is used to normalize the values.

Another type of surrogate model constructed at KLM is a Higher Dimensional Model Representation (HDMR). These can be used to create a model for a high dimensional system based on expert knowledge and data [5]. By decomposing high dimensional functions into subsets of lower dimensional component functions, the high dimensional functions can be approximated with high accuracy. Originally introduced by Sobel [38], it was proven that any function can be split into a set of sub-functions of different dimensions. Rather than treating a function as a black box, a HDMR allows a function to be represented in terms of the effects of individual variables and their interactions. This decomposition provides interpretability, which helps to identify the dominant factors and their combined influence on the system output [5].

A general HMDR system can be written as a constant term plus additional contributions from single variables and their interactions. This representation can be seen in Equation 4.6 [5]. A drawback of this system is its exponential growth: The number of terms increases exponentially with the number of

variables. For instance, a system with 5 variables requires 32 terms, while a system with 6 requires 64. To address this problem, the HMDR is usually truncated by ignoring terms above the second order (Second-Order HDMR). This increases computational speed drastically while not causing a big dent in accuracy. However, it does assume high-order terms are negligible, which may not hold for all systems.

$$f(\mathbf{x}) = g_0 + \sum_{i=1}^{m} g_i(x_i) + \sum_{1 \le i_1 < i_2 \le m} g_{i_1 i_2}(x_{i_1}, x_{i_2}) + \sum_{1 \le i_1 < i_2 < i_3 \le m} g_{i_1 i_2 i_3}(x_{i_1}, x_{i_2}, x_{i_3}) + \dots \tag{4.6}$$

4.3.2. Long-Short Term Memory

A Long-Short Term Memory (LSTM) neural network is a gated network to identify anomalies by calculating the difference between expected parameters and measured parameters [39]. The LSTM is designed to remember important patterns over long sequences, while forgetting irrelevant information. It does this using a forget gate, input gate and output gate. Each gate uses sigmoid-activated fully connected layers, like volume knobs that control information flow. For example, if it detects that during 45 flights the fuel flow stayed stable, the forget gate will decides it is not important so scales down its weight. If the input gate sees the HPT inlet temperature rising, it stores this trend in memory. The output gate decides this trend should affect the health, and the health indicator is changed.

4.3.3. Attention-Based Deep Learning

Chen et al. [40] proposed an Attention-Based Deep Learning method, where a LSTM neural network is combined with an attention mechanisms. While LSTM networks are designed to capture relevant information over long sequences, they are still susceptible to information loss, especially when critical events happen far from the prediction point. The integration of an attention mechanism mitigates this issue by enabling the network to dynamically assign weights to different time steps, thereby allowing the model to revisit and emphasize past observations that are most relevant for the prediction task.

An additional advantage of this approach is improved interpretability. By examining the attention weights, it is possible to identify which data input contributed most to the models output. This enhances transparency and provides valuable insights into the temporal dynamics of the engine's degradation.

4.3.4. Generative Adversarial Networks

Unsupervised learning is a broader form of machine learning aimed at discovering patterns or structures within datasets that lack explicit labels [41]. A Generative Adversarial Network is a subfield of unsupervised learning, where the goal is to train a model that approximates the underlying data distribution and generate new samples which are similar to the training data.

Generative Adversarial Networks (GANs) are one of the most widely used generative modeling techniques. They consist of two competing neural networks: a *generator* and a *discriminator*. The generator produces synthetic data with the objective of mimicking the real dataset, while the discriminator tries to distinguish between real and synthetic samples. The adversarial setup encourages the generator to improve until the discriminator can no longer reliably tell the difference between the two. An illustration of the GAN setup is shown in Figure 4.1.

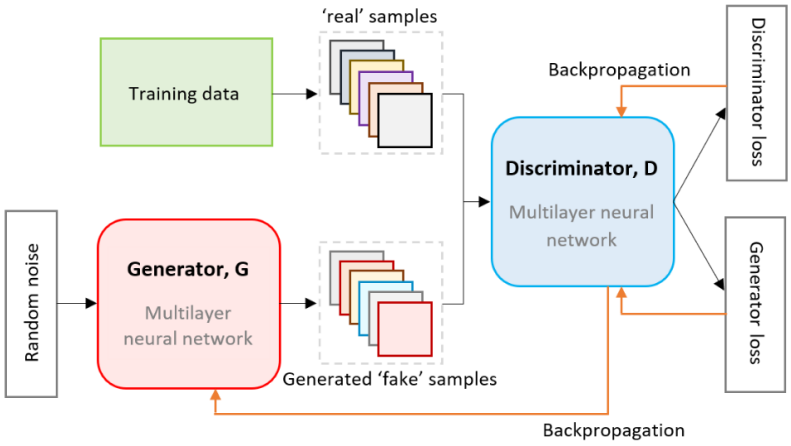

Figure 4.1: Example of a Generative Adversarial Network [42]

A GAN can be challenging to optimize, as it is hard to balance the training of the generator and discriminator. In addition, learning gradients may vanish when the discriminator does not give the generator enough information to learn. To address these challenges, several architectural improvements have been proposed. One of them is the Deep Convolutional GAN (DCGAN) [43], which replaces fully connected hidden layers with convolutional and de-convolutional layers. This makes the method strong for image data, but less effective for structured data.

4.3.5. Self Organising maps

Self-Organizing maps (SOMs) are a type of unsupervised neural network used for clustering, visualization and exploratory data analysis [44]. By projecting high-dimensional data in a low-dimensional 2D grid, it can process large datasets. Unlike supervised learning, it does not require labeling. Instead, the algorithm is based on competitive learning, where neurons compete to best represent the input vector. The neuron with the closest weight relative to the input neuron is chosen are the *winner* neuron, and its neighbors are updated using the Kohonen rule. Through this process, the map preserves relationships between cells, meaning that similar data points are mapped in cells close to each other. This makes SOMs particularly useful for exploring hidden patterns, grouping unlabeled datasets, and visualizing complex data structures. SOMs can therefore be used to cluster health indicators, and visualize complex engine data, which can be useful to the current project.

Advantages are that is can handle complex, high-dimensional data, and provide a visualization of that data in a map. Furthermore, a SOM does not require labels or other information prior to constructing the map. However, this does make it computationally expensive. Next to that, in order for the map to represent the data accurately, careful tuning is required to ensure a good learning rate and map size.

4.3.6. Differential Evolutionary Algorithm

To address the underdetermined nature of the GEnx-1B simulation problem, several optimization techniques have been explored. Among these, evolutionary algorithms have shown great promise due to their ability to iteratively search for near-optimal solutions in a multi-dimensional space. However, a traditional evolution algorithm operates on binary strings, which is not suited for the current situation. To overcome this limitation, a differential Evolution (DE) algorithm is used, which uses real-valued vectors instead of binary strings [3].

A DE is a population-based optimization approach that operates through a process of parallel direct search. The algorithm begins by initializing a number of parameter vectors (agents) which cover the entire search domain. New candidate solutions are then generated through a mutation process, where the difference between randomly selected vectors is scaled and added to a base vector. Crossover is the next step, where elements of the mutant vector are combined with the target vector resulting in a new trial vector. If the trial vector from the crossover operation has a lower objective function value than the target vector, it replaces the target vector in the next generation.

To employ this method, a substantial number of data points are required. The quality of the data is

vital as it directly influences the search space and convergence of the optimization. Rootliep et al. [3] implemented various filtering methods to ensure that the dataset was as accurate and noise-free as possible. This preprocessing step allowed the DE algorithm to accurately guide the search space toward optimal or near-optimal solutions. This method does employ GPA to run off-design simulations using health indicators as an input vector. Therefore this method can be seen as a hybrid-approach.

4.3.7. Rule-Based Methods

Another field of research involves rule-based methods, which correspond to traditional expert systems and fuzzy logic. These are similar methods, but differ in how they handle uncertainty and decision-making. Expert systems developed in the past relied on rule-based reasoning, like a chain of *IF-THEN-ELSE* rules. These rules are written in a knowledge database, which is setup by the developer. This lead to fast diagnostics and therefore quick maintenance actions could be performed [45]. In addition, since this system corresponds to human reasoning, it is very interpretable. However, due to the simplicity of the system, smearing is unavoidable and leads to misidentification. Another downside is that expert system output a binary value, stating if a failure has occurred or if a certain threshold has been exceeded [5].

Fuzzy systems were introduces to address these shortcomings. Unlike expert systems that operate with binary truths, fuzzy logic allows reasoning with degrees of certainty, enabling the system to handle noise and measurement errors [46, 23]. The heart of the fuzzy system is still the rule-base, just like expert systems, but also consists of a Input/Output pairs database. The I/O pairs can be extracted from the Fault Signature Table (FST). The FST, which is unique for each engine, specifies the relationship between deviations in health indicators and corresponding variations in measured parameters. This serves as the foundation of fuzzy systems.

4.4. Companion Engine Analysis

When applying MOPA to engine data, the analysis does not have to be restricted to one engine. By incorporating data from a fleet or by comparing paired *companion* engines on the same aircraft, additional data can be exploited to enhance diagnostics and prognostics. The central ideology of companion engine analysis is that companion engines operating under nearly identical operating conditions should exhibit the same gas path parameter data, if engines are of similar age and maintenance record, otherwise baselining is required. If companion engines do not exhibit similar gas parameter data, degradation or a failure of an engine component has occurred.

Companion engine analysis has been investigated in several studies as a way to enhance engine health monitoring by exploiting similarities between engines operating under comparable conditions. Singhvi [7] developed a framework at KLM where paired engines were compared through surrogate models and health matrices, with degradation indicators fed into an LSTM for fault classification. Verbist [27] explored a fleet-wide comparison approach using EGT margin but found that benchmarking against fleet averages offered limited diagnostic insight. Jacobs et al. [47] introduced an inter-engine variation analysis based on residuals filtered through a Mixture of Experts model, reducing sensitivity to noise while improving anomaly detection. Volponi et al. [37] proposed a cross-wing comparison method that combined state variable models with Kalman filtering, demonstrating strong fault isolation accuracy.

In practice, companion engine analysis can be implemented pairwise, or fleet-wise. Pairwise methods can allow detection of abnormal deterioration, while fleet-wise methods can compare individual engines to the fleet baseline. The main advantages are improved robustness to environmental variation, and the ability to detect abnormal degradation.

4.5. Conclusion

In general, the modeling approaches discussed will have to cope with the following requirements:

1. Can deal with a limited amount of measurements (Y < X)
2. Can deal with measurement noise and bias
3. Can find a solution within a reasonable time frame (high computational speed)
4. Can deal with non-linearity
5. Can avoid smearing

6. Can be interpreted

While a range of modeling approaches were investigated in this literature study, approaches that do not enable a direct prediction of health indicators from gas path parameters or provide a solution to the underdetermined problem will not be included in the comparison. The table with suitable modeling approaches is shown in Table 4.1.

Characteristic	GPA	GSP/ GSPy	Multi-Point GPA	GSPy-HIS	Residual Methods	Surrogate Models	HDMR	DEA	Expert Systems	Fuzzy Logic
Can deal with Y<X	No	No	Yes	Yes	Yes	Yes	Yes	Yes	No	No
Can deal with measurement noise	No	No	No	No	No	No	No	No	No	Yes
High computational speed	Yes	Yes	Yes	Yes	Yes	Yes	No	No	Yes	Yes
Can deal with non-linearity	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No
Can avoid smearing	No	Yes	Yes	Yes	No	No	Yes	Yes	No	Yes
Can be interpreted	Yes	Yes	Yes	Yes	Yes	No	Yes	No	Yes	Yes
References	[12]	[29, 8]	[34, 3]	[35, 36]	[7]	[26]	[5]	[3]	[45, 5]	[46, 23]

Table 4.1: Trade-off of modeling approaches; GSPy-HIS: GSPy Health Indicator Selection; DEA: Differential Evolutionary Algorithm

The most suitable approach is the use of GSPy with health indicator selection for the prediction of engine health indicators. This approach satisfies the key requirement of handling the underdetermined system, while also being able to capture non-linear behavior and maintain interpretability. Although GSPy has limitations in dealing with measurement noise, this can be addressed by implementing measurement tolerances within the model. A key advantage of GSPy is its flexibility as an open and adaptable framework, allowing the model to be modified to meet specific modeling requirements.

Companion Engine Analysis will also be used as a modeling approach, although it is applied after the prediction of health indicators. By comparing trends between companion engines, deviations can be identified, enabling the detection of relatively more degraded engines. This relative comparison provides an extra source of information and can reduce dependence on model accuracy. As a result, it can give additional information not visible from absolute health indicator trends.

Part II

Methodology: Component Exclusion Method

5

Implementation Platform

Building on the findings of the literature review presented in Part I, the methodology is developed in Part II. GSP and GSPy have been introduced in brief earlier, a more comprehensive explanation will now be provided. In addition, adaptations made to GSPy to improve the match between GSPy, GSP, and on-wing data will be introduced. Finally, verification is done on the GSPy model.

5.1. GSPy

GSPy is a Python-based gas turbine simulation tool inspired by GSP, the Gas Turbine Simulation Program. It is open-source and available on GitHub [8]. The software replicates GSP's component-based architecture using object-oriented programming, in which individual objects encapsulate both data and functions and interact with each other. In this framework, each gas turbine component is represented as an object that receives thermodynamic states as input, applies the appropriate physical models, and returns an updated state to the subsequent component. The overall layout of a two-spool turbofan engine, such as the GEnx-1B, is illustrated in Figure 5.1.

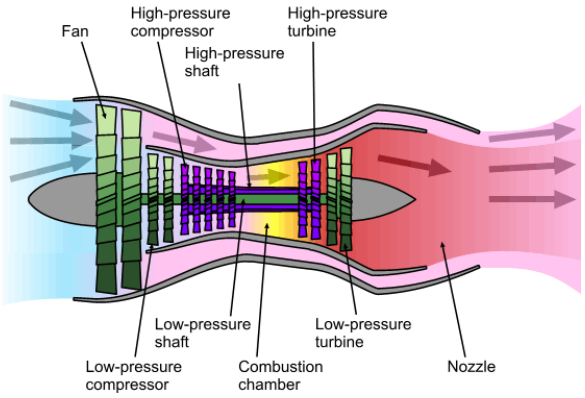

Figure 5.1: Schematic of two-spool turbofan engine [48]

GSPy supports both design point and off-design simulations across varying operating conditions. For turbomachinery components, compressor and turbine maps are used to determine pressure ratio, efficiency, corrected mass flow, and rotational speed. Thermodynamic properties are computed using the open-source Cantera library, enabling the modeling of real-gas mixtures and combustion processes. A numerical solver is used to solve the resulting non-linear system of residual equations.

In GSPy, a model is constructed by specifying the components in the correct order according to their station number. For the specific GEnx model in which the Fan core and LPC are combined into one component the sequence is the following: Inlet - Fan - HPC - Combustor - HPT - LPT - Exhaust. Each component may require associated characteristic maps (e.g. compressor or turbine maps) to

define its off-design behavior. The ordering of components is critical, as each component receives the thermodynamic output of its predecessor as input. During simulation, GSPy passes a gas state object through the component chain; each component modifies the state through compression, expansion, or heat addition, and the updated state is then forwarded to the next component in the sequence. This can be seen in Figure 5.2, where the arrows represent the flow of information.

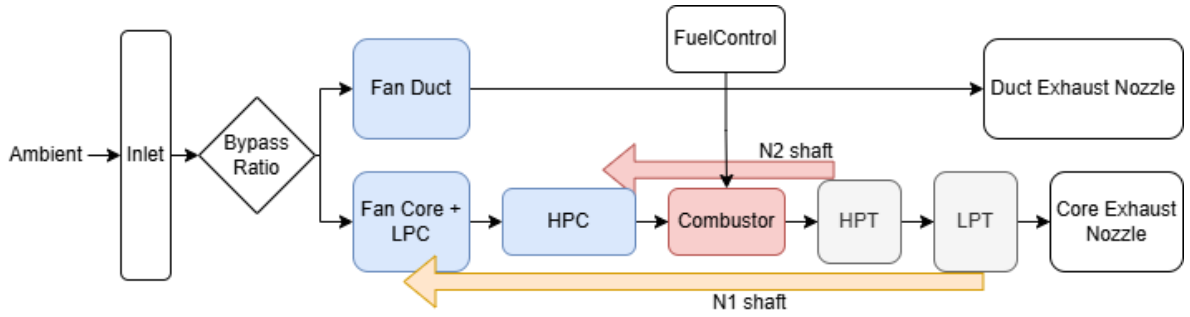

Figure 5.2: GSPy GEnx-1B component pipeline

A more detailed description of the GSPy simulation framework, including the component structure and simulation steps is provided in Appendix B.

5.2. GSPy Customizations & Extensions

This section outlines modifications introduced to GSPy to improve its applicability to engine performance and monitoring, specifically targeting the GEnx-1B model and its engine data. While the baseline GSPy code supported Adaptive Modeling calculations for a generic turbojet engine, additional adaptations were necessary to achieve a robust and accurate representation of the GEnx-1B. These updates can be split into two main aspects: Adaptation of Adaptive Modeling module by implementing N1 control, measurement tolerances and health indicator bounds; and implementing crossflow calculations, ensuring a better model match between GSPy and GSP, and more flexibility to match the GSPy model to on-wing data.

5.2.1. Simulated Deterioration

Initially, the Adaptive Modeling (AM) capabilities of GSPy were evaluated using simulated deterioration generated in GSP. However, discrepancies between GSP and GSPy led to inconsistent results, making it difficult to distinguish modeling errors in GSPy from differences between the two simulation environments.

To address this, a dedicated class was developed to model simulated deterioration directly within GSPy. This approach applies scaling factors to selected health indicators at each run point, enabling controlled deterioration to be introduced within an off-design simulation. As a result, the performance of the AM module can be assessed more reliably, without the influence of discrepancies between GSP and GSPy.

5.2.2. Adaptive Modeling

Before applying the Adaptive Modeling module for health indicator prediction using on-wing data, several features were implemented to enhance its functionality and accuracy.

N1 Control

The AM module is GSPy was originally designed to have the fuel flow as an input. However, in flight the pilots set N1 and a computer adapts the fuel flow as such. Therefore implementing the N1 controller to work improves the compatibility with on-wing data. The N1 controller is implemented in the following way. In the `TAMcontrol` class a check was created to find what the power setting is. If this is not fuel flow or None, for example N1, an additional error equation is added:

$$r_{P_{set}} = \frac{P_{set,meas} - P_{set,model}}{P_{set,DP}} \tag{5.1}$$

This is the same method as the `Control` class to set N1 control in a simulation.

Measurement Tolerances

Table 2.3 shows that the sensors are subject to uncertainty. To address this, tolerances are introduced to the Adaptive Modeling module, allowing the simulation to accommodate measurement uncertainty. As described in subsection B.2.4, the AM module introduces residual equations that compare each measured value to the corresponding simulation output. To account for sensor uncertainty, a tolerance factor is applied to each residual. This approach ensures that the model can still converge even when the measurements are slightly offset from their true values. The tolerance for each health indicator is incorporated into the residual calculation as shown in Equation 5.2:

$$r_i = T_i \cdot \frac{\text{Output}_i - \text{Measured}_i}{\text{Model}_{\text{des,i}}} \tag{5.2}$$

Here, r_i is the residual for health indicator i, T_i is the selected tolerance, Output_i is the simulation output, Measured_i is the measured value, and $\text{Model}_{\text{des,i}}$ is the model's design value for normalization.

For the GEnx-1B case, the AM function uses a system error tolerance of 0.001. Therefore, the residuals must be scaled according to the expected measurement uncertainty of each sensor. For example, if a sensor has a measurement uncertainty of 1%, a tolerance factor of 0.1 is applied so that the resulting residual magnitude remains comparable to the AM system error tolerance. Without this scaling, measurement uncertainties can dominate the residual calculation, causing many Adaptive Modeling simulations to fail to converge.

Health Indicator Bounds

Finally, health indicator bounds were introduced to the AM module. Initial testing of the AM approach with on-wing data revealed that, although the simulation successfully converged, it occasionally produced unrealistic values for certain health indicators. To mitigate this issue, explicit bounds were imposed on health indicators during the simulation process.

The root solver does not allow for bounds as an input of the function. Therefore, an alternative strategy was developed to enforce these bounds. Health indicator values are treated as system states within the AM module. During each iteration, the AM `PostRun` function compares the state value of each health indicator to its bound. If exceeded, a penalty term is applied: The squared difference between the state value and the bound, scaled by a designated factor, is added to all residuals. This approach discourages the solver to converge to solutions outside the acceptable range.

A disadvantage of this approach is the additional computational time required by the solver when the optimal solution lies outside the bounds. The solver might spend considerable time searching for a feasible solution within the bounded domain, leading to a high number of iterations. This can lead to non-convergence errors if no acceptable solution is found. Therefore, a different implementation of bounds was used, explained in section 7.3.

5.2.3. Crossflow Implementation

Vital verification of GSPy was performed using a single-spool turbojet model, for which deviations with respect to GSP were below 0.2%. However, when applying GSPy to a high-bypass turbofan model of the GEnx engine, developed at KLM Engine Services, deviations exceeding 3% were observed across multiple operating conditions. These discrepancies indicated that the turbofan flow split modeling in GSPy was not fully consistent with GSP. A detailed comparison revealed a difference in the treatment of the core and bypass mass flow split at off-design conditions. In GSP, the split between core and bypass flow is governed by the bypass ratio influence factor, denoted by Cf, which accounts for crossflow induced by deviations of the bypass ratio from its design value.

As the flow enters the fan, it is divided into a core stream and a bypass stream. At the design point, the bypass ratio BPRdes defines this split. Under off-design conditions, however, the actual bypass ratio BPR generally differs from the design value, which induces a redistribution of mass flow between the core and bypass ducts. GSP accounts for this effect by introducing a crossflow term, defined as the difference between the duct mass flow corresponding to the off-design bypass ratio and that corresponding to the design bypass ratio [49]. This crossflow is given by Equation 5.3.

$$\Delta\dot{m}_{dBPR} = \dot{m}_{in} \cdot \left(\frac{BPR}{BPR+1} - \frac{BPR_{des}}{BPR_{des}+1} \right) \tag{5.3}$$

Where $\dot{m}_{in}$ is the total inlet mass flow. The bypass influence factor Cf can be interpreted as the position of the dividing streamline. It is used to calculate the core and duct mass flow represented in Equation 5.4 and Equation 5.5.

$$\dot{m}_{\text{duct}} = \dot{m}_{\text{in}} \cdot \frac{BPR_{\text{des}}}{BPR_{\text{des}}+1} + Cf \cdot \Delta\dot{m}_{\text{dBPR}} \tag{5.4}$$

$$\dot{m}_{\text{core}} = \dot{m}_{\text{in}} \cdot \frac{1}{BPR_{\text{des}}+1} - Cf \cdot \Delta\dot{m}_{\text{dBPR}} \tag{5.5}$$

The factor Cf is bounded between 0 and 1 and represents the degree to which off-design bypass ratio variations influence the fan compression process. Two limiting cases can be identified.

In the first limiting case, $Cf = 0$, the off-design bypass ratio has no influence on the fan compression process, as illustrated in Figure 5.3a. In this scenario, both the core and bypass flows follow the design-point dividing streamline through the fan. The fan maps are therefore evaluated using the design bypass ratio and the total inlet mass flow, and the redistribution to the off-design bypass ratio occurs only downstream of the fan, at the splitter.

In the second limiting case, $Cf = 1$, shown in Figure 5.3b, the dividing streamline remains fixed at its design-point position. Consequently, the off-design bypass ratio fully affects the fan compression process, and both the inlet mass flow and the off-design bypass ratio directly determine the operating points on the core and bypass fan maps.

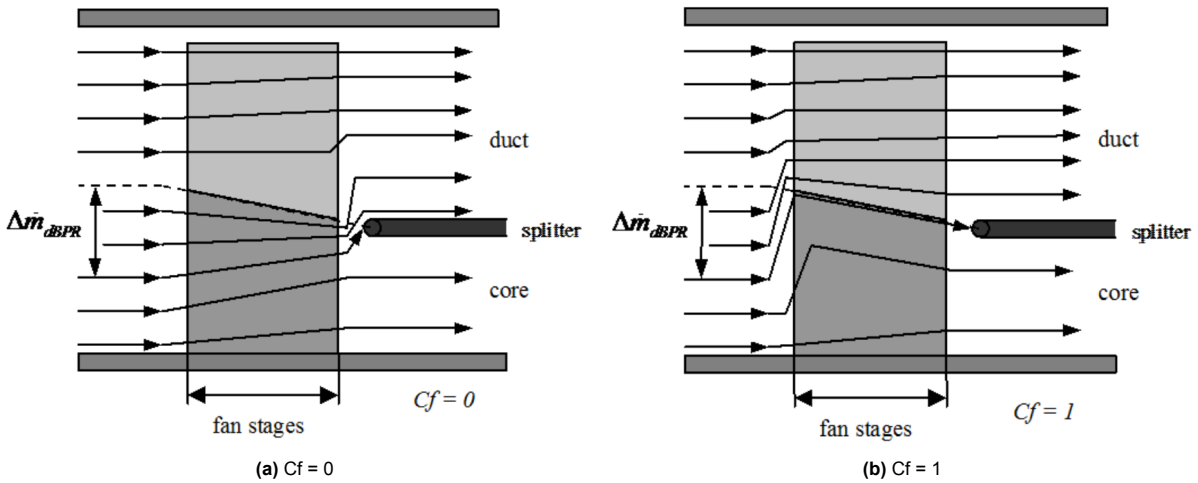

(a) Cf = 0 **(b)** Cf = 1

Figure 5.3: Bypass ratio influence factor Cf visualization [49]

By introducing the bypass ratio influence factor into GSPy, the mass flow split at off-design conditions was made consistent with the formulation used in GSP, resolving a key source of discrepancy for high-bypass turbofan configurations. Furthermore, it is uncertain what the on-wing Cf factor is, this can be tuned during validation. To verify the correct implementation of the bypass ratio influence factor, a systematic comparison between GSPy and GSP was performed. The resulting bypass ratio variation as a function of fan spool speed is shown in Figure 5.4.

Several observations can be made based on this comparison. Most notably, for the case $Cf = 0$ in GSPy, the bypass ratio decreases significantly at operating points with fan speeds exceeding the design point. This behavior can be attributed to differences in the thermodynamic property evaluation between the two tools. In GSPy, the gas properties of the crossflow are mixing with the duct/core stream differently than in GSP, causing these discrepancies.

In addition, an overall trend is noticeable: The higher the spool speed, the lower the bypass ratio. This trend is consistent with engine operation, as an increased spool speed creates a suction effect on the core stream, increasing core mass flow more rapidly than bypass mass flow.

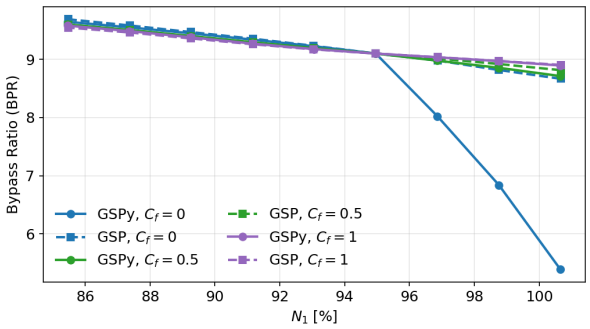

Figure 5.4: BPR vs N1 spool speed comparison of GSP and GSPy with different Cf factors

Finally, a small trend between the Cf curves is visible: At higher Cf factors, the slope of the bypass ratio is lower. This is because the OD bypass ratio does not play a role in the maps compression at low Cf factors.

Comparing the temperature and pressure ratio's over the fan (Figure 5.5) the most noticeable is the effect of the Cf factors. A higher Cf factor causes a *flatter* curve. This means changes in the corrected mass flow result in a smaller change in temperature and pressure ratio. In addition, the curves all match below the design-point, while the curve of Cf=0.5 and Cf=1.0 do not match above the design point.

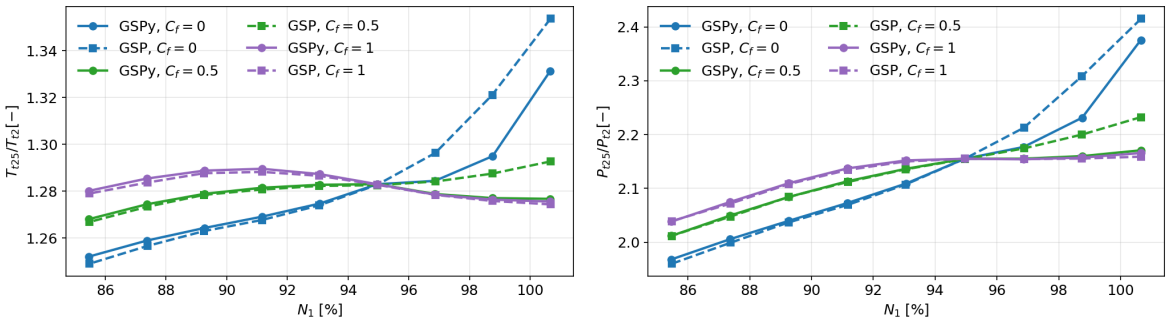

(a) Temperature ratio over the fan vs N1 spool speed

(b) Pressure ratio over the fan vs N1 spool speed

Figure 5.5: Temperature ratio and pressure ratio comparison of GSP and GSPy with different Cf factors

5.3. Model Verification

Since GSPy is a new software, verification has to be performed to ensure it gives accurate solutions. Initially, the assumption was made that simulations in GSP using the GEnx-1B model constructed by KLM was correct. This was then used to verify GSPy. Since at first simulations of the GEnx-1B model in GSPy and GSP did not match, additional verification on the turbojet model was done. In this section the verification of the turbojet model and the GEnx-1B model will be explained.

5.3.1. Turbojet Model

Verification of the turbojet was primarily performed by Visser prior to publishing GSPy. However, this was expanded upon by verifying another turbojet configuration. Two configurations were analyzed: **Model A**, representing the original verification case, and **Model B**, an additional configuration representing the sample model in GSP. The objective of the verification was to assess consistency between GSP and GSPy.

Model A is the same as the turbojet model in GSPy. This model was run and matched with the verification results of Visser. Deviations remain below 0.2% for all engine parameters. At off-design conditions further from the design point, deviations increase slightly but remain small.

Model B was evaluated using the GSP sample model and an equivalent model made in GSPy. The adjusted parameters compared to **Model A** were the following: The combustor efficiency, the combustor

pressure loss and the nozzle velocity coefficient.

Compared to model A, slightly larger deviations were observed. Around the design point, deviations are small (<0.2%) except for the net thrust, which has a 0.5% deviation. This discrepancy is due to the nozzle velocity coefficient (CV). The implementation of the CV differs between GSP and GSPy. To conclude, both models show matching simulation results in GSP and GSPy. This means for the turbojet model, GSPy is able to match GSP.

5.3.2. GEnx-1B Model

A GSPy turbofan model was constructed by replicating the design-point values from a GSP GEnx model developed at KLM [33]. Steady-state series simulations for N1 spool speeds were then performed. This was done in two operating conditions: Design point (test-cell) and off-design (cruise).

Initially, large deviations between the GSP and GSPy were observed. Since the turbojet model was verified, the source of the error was found to be the calculation and usage of the bypass ratio in the fan. The crossflow was not yet implemented, and this let to major discrepancies. In the GSP GEnx-1B model, the Cf factor was set at zero, while in GSPy, since the Cf factor was not implemented, it was effectively at 1.

Design-Point Operating Condition

Since the model is creating using test-cell data, the verification simulations very first performed at these conditions. To correctly align the models, the required isentropic efficiencies were obtained from the thesis of Ramdin [33]. Using these values, new GSP and GSPy models were constructed.

The outcomes of the verification simulations are presented in Figure 5.6. A Cf factor of 0.5 is used in all simulations. The N1 speed was included as a reference metric to verify the overall consistency of the simulations. In the legend, the corrected N1 spool speed of all simulations is shown, as this parameter provides a reliable basis for comparison because it corrects for ambient temperature.

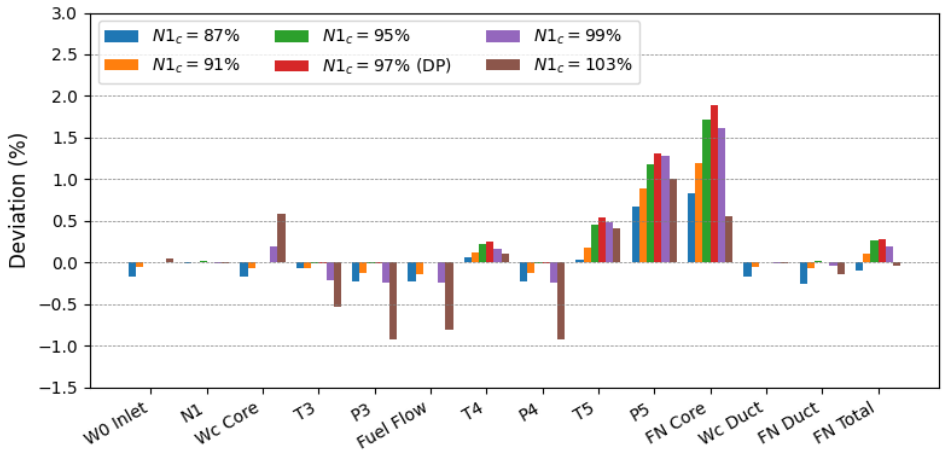

Figure 5.6: Differences between GSP and GSPy for the GEnx model using an off-design st.st. simulation with N1 spool speed as control variable; Positive deviation = GSPy value > GSP value

At the design-point, a discrepancy is already observed downstream of the combustor. This deviation originates from the approaches used by GSP and GSPy to model heat release. In GSP, a higher amount of chemical dissociation is predicted, which effectively lowers the heat release and leads to lower post-combustion temperatures. After careful examination, it was found that GSPy employs a slightly more advanced chemical equilibrium calculation. Therefore, the thermodynamic variables calculated by GSPy post-combustion are considered more accurate than those obtained from GSP.

For off-design (OD) simulations at corrected fan speeds below the design point, the alignment between GSP and GSPy remains very good. However, when operating above the design-point corrected speed, the OD simulation shows larger deviations. This comes from the crossover flow calculation. If the OD N1 speed is higher than the DP N1 speed, some of the duct flow will be sucked into the core. The effective bypass ratio is different than the design bypass ratio. Due to this suction, the core flow will be

enriched with some of the bypass flow. Consequently, the flow properties of the core flow will change. Due to the slightly difference chemical calculation in GSPy and GSP, this will give small deviations, which then propagate through the core engine, leading to deviations in downstream temperatures and pressures.

Off-Design Operating Conditions

To further verify the performance of GSPy under realistic flight conditions, additional simulations were conducted for cruise operation. Both GSP and GSPy were run at the following off-design conditions: An altitude of 10,000 m, and a Mach number of 0.8. The resulting deviations between GSP and GSPy are shown in Figure 5.7.

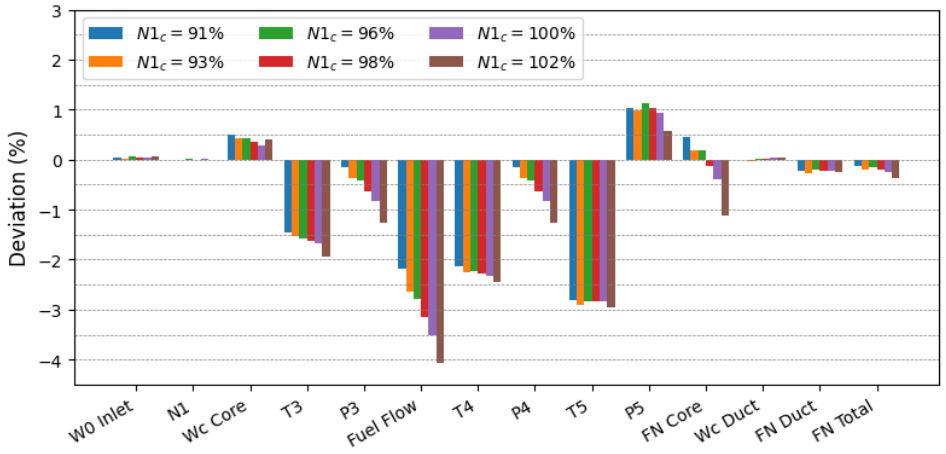

Figure 5.7: Differences between GSP and GSPy for the GEnx model in cruise conditions using an off-design st.st. simulation with N1 spool speed as control variable; Positive deviation = GSPy value > GSP value

As expected, the discrepancies between GSP and GSPy are slightly larger under cruise conditions than at the design point. The most pronounced differences are observed in the fuel flow and temperature-related variables, which can again be attributed to differences in the underlying chemical equilibrium and heat release modeling. Furthermore, the corrected core mass flow is higher in GSPy, and to match the input N1 spool speed, the fuel flow is lower.

5.4. Adaptive Modeling Verification

Adaptive Modeling (AM) is a critical component in the calculation of health indicators and thus its verification is essential. First, adaptive modeling was done using the turbojet model and its consistency between GSP and GSPy verified. Then, this was also done for the GEnx-1B model.

5.4.1. Turbojet AM Verification

For the turbojet verification, Model A was used. The verification process begins with loading the TJET model in GSP and creating a steady-state series case. This is used to systematically vary the compressor efficiency, compressor mass flow, turbine efficiency, and turbine mass flow. Since the efficiency and mass flow are correlated [3], a 1% efficiency deterioration is coupled to a 2% mass flow deterioration. For the turbine, the mass flow deterioration is a positive increase in mass flow capacity. The outputs from these runs, including temperatures, pressures, and fuel flow, are saved and subsequently incorporated into a log file for the AM module.

For GSPy, the AM module is pre-configured for the turbojet, but an input file must be generated from the GSP simulation outputs. Once the input file is prepared, the AM module in GSPy is executed, and the resulting outputs are recorded for comparison with GSP outputs. The goal of this verification is to ensure that the AM module reproduces the inputs accurately under controlled degradation conditions.

The verification relies on a set of key measurements, specifically the total temperature and total pressure at the compressor exit, the total temperature and total pressure at the turbine exit, and the fuel flow. While five measurements were available, only four are strictly necessary for the AM calculations,

with fuel flow always retained as a reference input. Analysis of different measurement subsets revealed that the compressor exit pressure is essential for convergence. Excluding this measurement prevents the AM simulation from converging.

Moreover, using only the compressor exit pressure without the compressor exit temperature leads to significant inaccuracies, with zero input degradation producing up to 22% deviation. Therefore, both compressor exit temperature and pressure must be included. While inclusion of turbine exit pressure slightly improves results compared to turbine exit temperature, the latter was chosen for reporting, as it aligns with the sensor configuration of the GEnx engine and represents a conservative, worst-case scenario.

The results obtained from the Adaptive Modeling (AM) tool in both GSP and GSPy at the design point (DP) are presented in Figure 5.8. As shown in the figure, both GSP and GSPy are capable of accurately predicting system deterioration. While GSP exhibits slightly higher fidelity, the performance of GSPy remains closely comparable, demonstrating its effectiveness even at the DP.

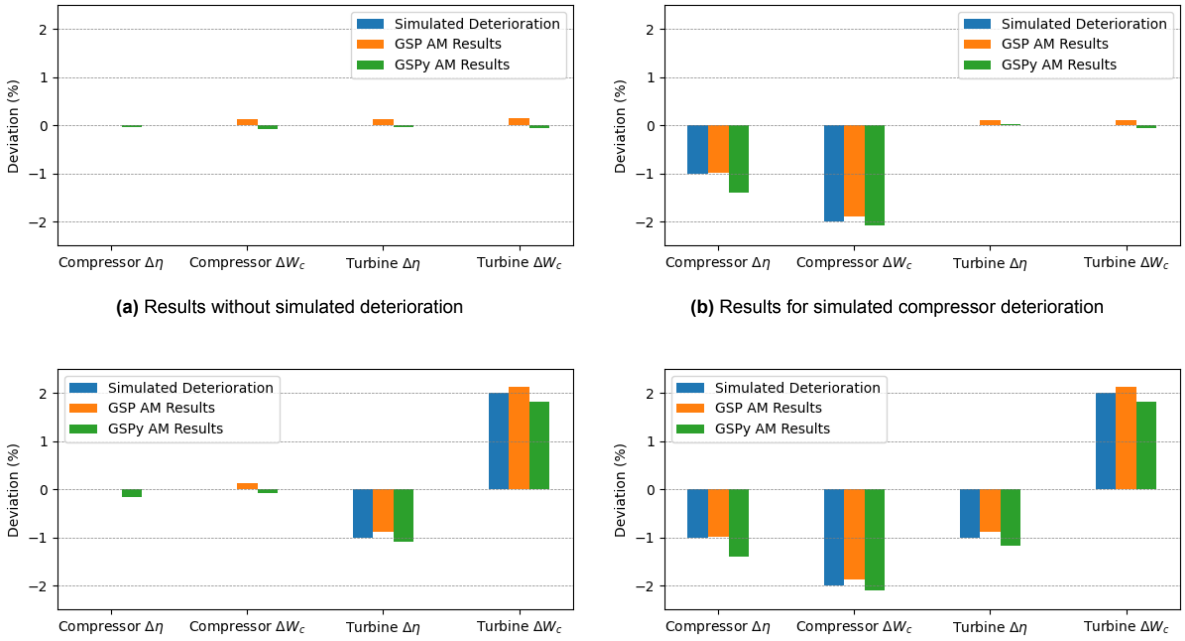

(a) Results without simulated deterioration

(b) Results for simulated compressor deterioration

(c) Results for simulated turbine deterioration

(d) Results for simulation compressor and turbine deterioration

Figure 5.8: Comparison of GSP and GSPy AM simulations for the turbojet model at DP conditions, with simulated deterioration

This adaptive modeling framework for the turbojet was further verified under off-design operating conditions:

1. **Cruise Condition 1:** The aircraft is simulated at an altitude of 5,000 m with a Mach number of 0.4. Ambient pressure and temperature are assumed to follow standard ISA conditions, and the corrected N1% is maintained at 100%.
2. **Cruise Condition 2:** The aircraft is simulated at an altitude of 10,000 m with a Mach number of 0.8. Ambient conditions are again set according to ISA standards, and the corrected N1% remains at 100%.

The simulation outputs obtained under these off-design conditions were stored and subsequently used as inputs for the adaptive modeling (AM) module in both GSP and GSPy. The corresponding results are presented in the following figure (Figure 5.9).

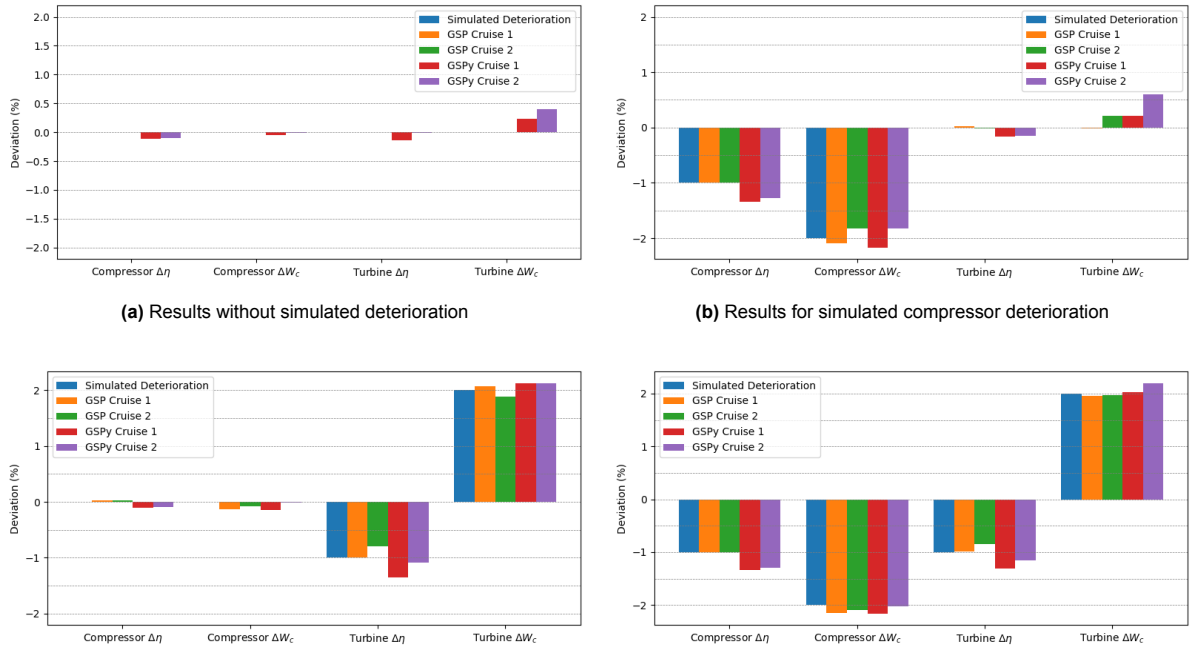

(a) Results without simulated deterioration

(b) Results for simulated compressor deterioration

(c) Results for simulated turbine deterioration

(d) Results for simulation compressor and turbine deterioration

Figure 5.9: Comparison of GSP and GSPy AM simulations for the turbojet model at OD conditions, with simulated deterioration

The Adaptive Modeling (AM) tool for the turbojet engine has been successfully verified within GSPy. Both design point and off-design simulations demonstrate that GSPy can reliably reproduce system performance and predict deterioration trends.

5.4.2. GEnx-1B Model AM Verification

The GEnx-1B model differs substantially from the turbojet model considered previously, both in architecture and in the availability of measurement data. In particular, the on-wing dataset for the GEnx contains only a limited number of sensors, which constrains the set of health parameters that can be reliably estimated. The verification of AM with the GEnx-1B model will be performed using a complete synthetic sensor set, to confirm that the methodology performs correctly when sufficient information is available. In the next chapter verification will be done using the on-wing sensor set. To fully remove inconsistencies in the adaptive modeling verification, a method was constructed to simulate deterioration scenarios in GSPy. This effectively removed the modeling offset created by differences between GSP and GSPy.

In order to create input degradation accurately representing real life deterioration, a set of input deterioration scenarios was created. These are found in Table 5.1, where η denotes efficiency and W_c denotes flow capacity. Each test mode will be simulated using an OD simulation in GSPy or GSP, and the gas path parameters will be stored. These values are then used for adaptive modeling. To assess the achievable performance, the GEnx AM model was tested using the following set of measurements:

- Temperatures: T_{25}, T_3, T_{49}
- Pressures: P_{25}, P_{s3}, P_{49}
- Rotor Speed: $N_2\%$
- Fuel Flow: W_f

In addition, a bypass ratio influence factor *Cf* of 0.5 was chosen.

Test Modes	Fan $\Delta\eta$	Fan ΔW_c	HPC $\Delta\eta$	HPC ΔW_c	HPT $\Delta\eta$	HPT ΔW_c	LPT $\Delta\eta$	LPT ΔW_c
Fan	-1	-2	0	0	0	0	0	0
HPC	0	0	-1	-2	0	0	0	0
HPT	0	0	0	0	-1	2	0	0
LPT	0	0	0	0	0	0	-1	2
Core	0	0	-1	-2	-1	2	0	0
Full	-1	-2	-1	-2	-1	2	-1	2

Table 5.1: Adaptive modeling verification test sets

The results of the DP AM simulation can be seen in Figure 5.10. Both GSP and GSPy are able to predict simulated deterioration to 0.1% accuracy. This was repeated for cruise conditions. The results were similar to DP results.

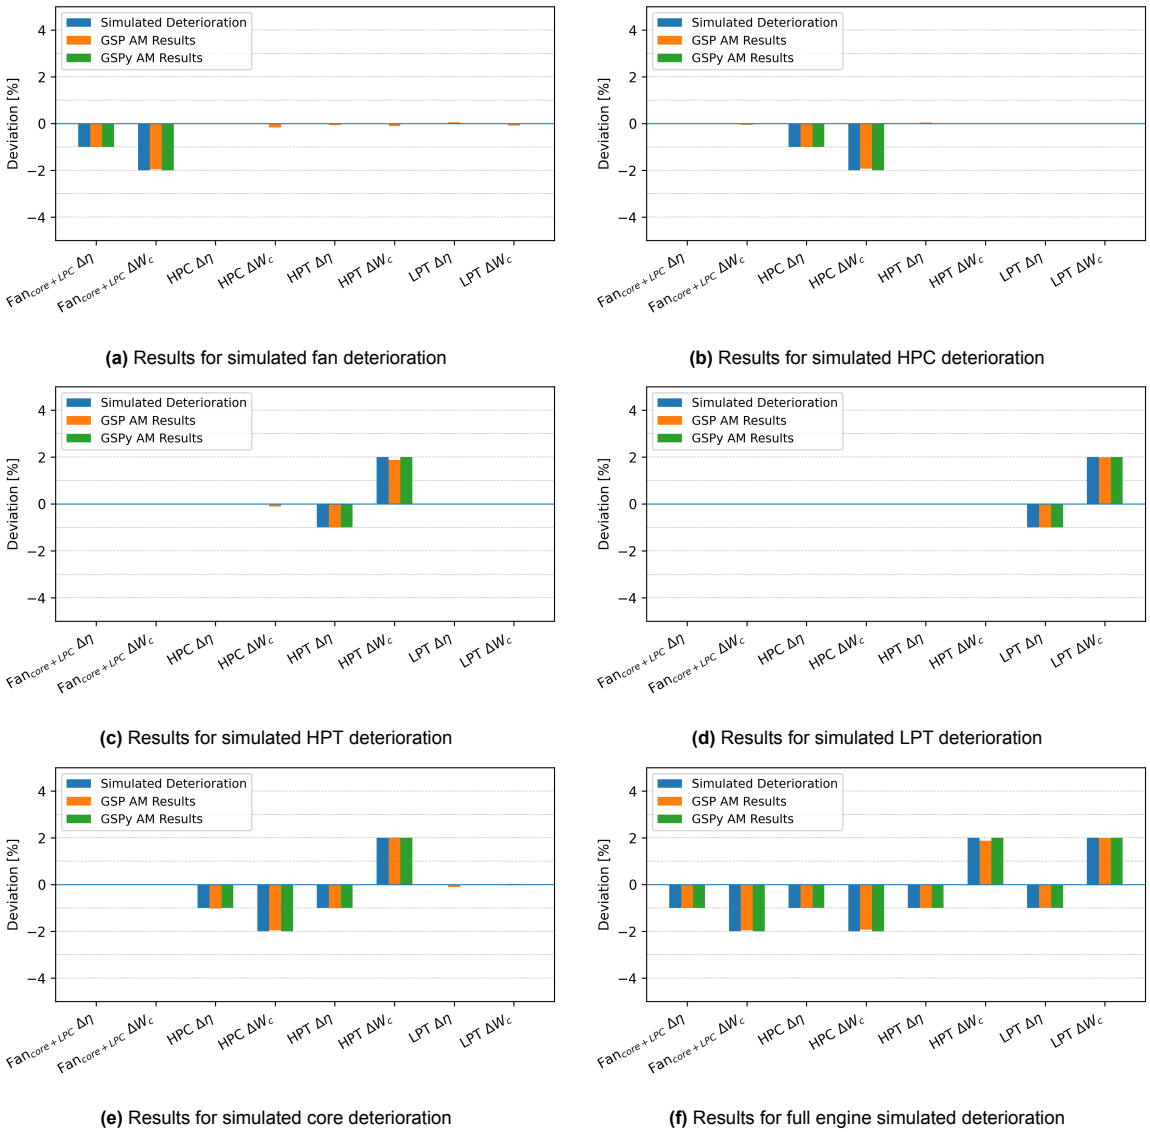

(a) Results for simulated fan deterioration

(b) Results for simulated HPC deterioration

(c) Results for simulated HPT deterioration

(d) Results for simulated LPT deterioration

(e) Results for simulated core deterioration

(f) Results for full engine simulated deterioration

Figure 5.10: Comparison of GSP and GSPy AM simulations for the GEnx model at DP conditions, with simulated deterioration

The practical challenge lies in applying the AM procedure to the on-wing GEnx dataset, which contains only six measurements: $T_{t,25}$, $T_{t,3}$, $P_{s,3}$, $T_{t,49}$, W_f, N_2. This limited set cannot fully constrain all eight

health parameters simultaneously, a systematic evaluation was conducted to determine which subsets of component indicators yield stable and accurate parameter estimates. Subsets were constructed where each subset contained 6 out of 8 possible health indicators. This gives a total of 28 subsets. In the next chapter the subsets will be used.

5.5. Conclusion

In this chapter, the GSPy framework has been successfully implemented, and verified to apply to turbofan engine performance analysis and condition monitoring. Key modification, such as the implementation of N1 control, measurement tolerances, health indicator bounds, and crossflow modeling, improved the consistency between GSP and GSPy, and improve model matching between the GSPy model and on-wing engines.

Verification results demonstrate that GSPy is capable of accurately reproducing GSP simulations at both design-point and off-design operating conditions, with small discrepancies being attributed to differences in thermodynamic modeling. Furthermore, the Adaptive Modeling (AM) module was thoroughly verified using both turbojet and GEnx-1B configurations, showing that GSPy can reliably estimate health indicator variations under controlled deterioration scenarios.

6

Gas Path Analysis on the GEnx-1B

This chapter will present a method to performing GPA on the GEnx-1B. This method consists of using subsets of health indicators to create a determined system. After an introduction into the subsets, the performance of each subset is tested and illustrated. Finally, a conclusion following the performance tests is presented.

6.1. Health Indicator Subsets

Adaptive Modeling (AM) requires that the number of condition parameters is equal to or greater than the number of measured performance variables to find an unique solution. This means that per simulation, six health parameters will be selected, forming a subset. Within a given subset, the selected health indicators act as free variables in the system of equations, allowing the solver to adjust them in order to achieve convergence. Health indicators that are not included in the subset remain fixed and are therefore not modified during the solution process.

In this study, the health indicators consist of the efficiency and flow capacity of the Fan core + LPC, HPC, HPT, and LPT. These components yield a total of eight health indicators. All possible combinations of six health indicators selected from the total of eight are considered, with each combination defined as a subset. In total, this results in a combinatorial number of subsets given by the binomial coefficient $\binom{8}{6} = \frac{8!}{6!2!} = 28$, based on factorial notation. These subsets can be applied in adaptive modeling for both simulated deterioration scenarios and on-wing sensor data. The subsets are illustrated in Figure 6.1, where the white spaces represent health indicators not included in that subset.

During the evaluation of all subsets, it was observed that certain combinations fail to converge under specific simulated deterioration conditions. These subsets are classified as ill-conditioned and excluded from further analysis. These subsets are colored in red. In Appendix C an analysis is provided on the non-converging subsets. The remaining twelve converging subsets are retained and form the basis for the subsequent analyses presented in this thesis.

6.2. Diagnostic Performance of Subsets

These 12 converging subsets were evaluated using seven simulated deterioration cases shown in Figure 6.2. The first case represents a healthy engine condition and serves as a baseline to verify the consistency and correctness of the adaptive modeling implementation. The performance of each subset was investigated for each deterioration case in terms of their ability to predict the imposed component deterioration.

Figure 6.1: CEM subsets

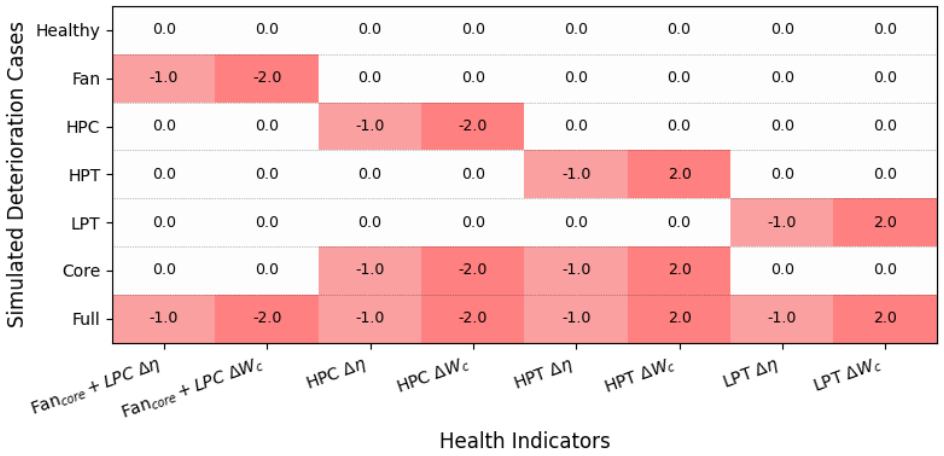

Figure 6.2: Simulated Deterioration Cases: Each value representing the percentage change of a certain health indicator

6.2.1. Fan Simulated Deterioration

The first deterioration case is the fan core + LPC deterioration, where the efficiency is reduced by 1% and the flow capacity by 2%. Subsets 1 through 6 include both the fan core + LPC efficiency and flow capacity health indicators, enabling accurate estimation of the imposed deterioration. The results demonstrate that when both efficiency and flow capacity are included as free parameters, the adaptive model successfully captures the degradation. The results of the first two of these six subsets for fan core + LPC simulated deterioration are seen in Figure 6.3, where the red labels are the health indicators not included in that subset.

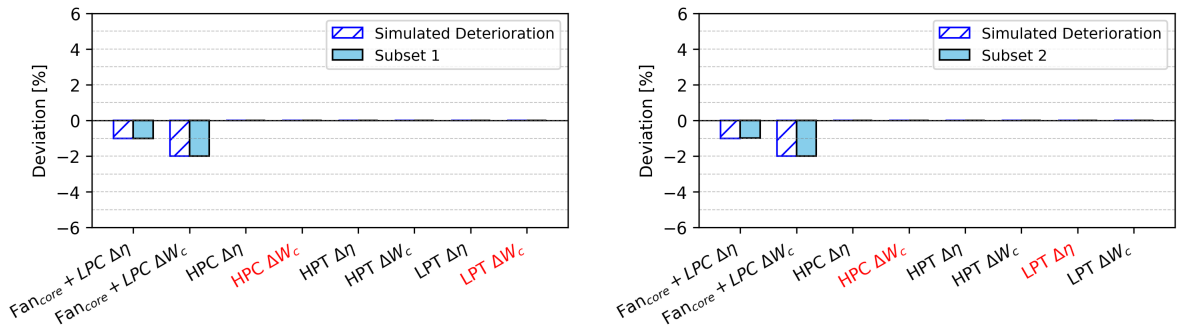

(a) Subset 1 **(b)** Subset 2

Figure 6.3: Subsets 1 and 2 with simulated fan core + LPC deterioration

In contrast, if the fan core + LPC efficiency or flow capacity health indicator is not included in a subset, the accuracy of health indicator predictions is significantly reduced. This can be seen in the last six subsets, of which two are shown in Figure 6.4. Subsets 7, 8 and 9 do not include the fan core + LPC efficiency, while subsets 10, 11 and 12 do not include the fan core + LPC flow capacity. If one of these is not included, the fan deterioration is redistributed to other components, a phenomenon called smearing. This effect is particularly pronounced when the flow capacity is excluded, leading to substantial and physically unrealistic deviations in other components.

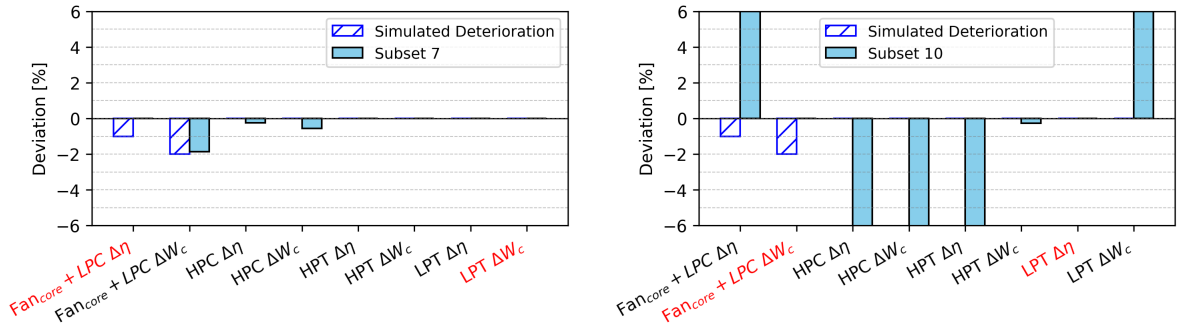

(a) Subset 7 **(b)** Subset 10

Figure 6.4: Subsets 7 and 10 with simulated fan core + LPC deterioration

6.2.2. HPC Simulated Deterioration

For the HPC deterioration case, subsets 7 through 12 include both HPC efficiency and flow capacity as health indicators, resulting in accurate predictions of the imposed degradation (Figure 6.5). These subsets demonstrate that inclusion of all relevant health indicators is essential for reliable predictions.

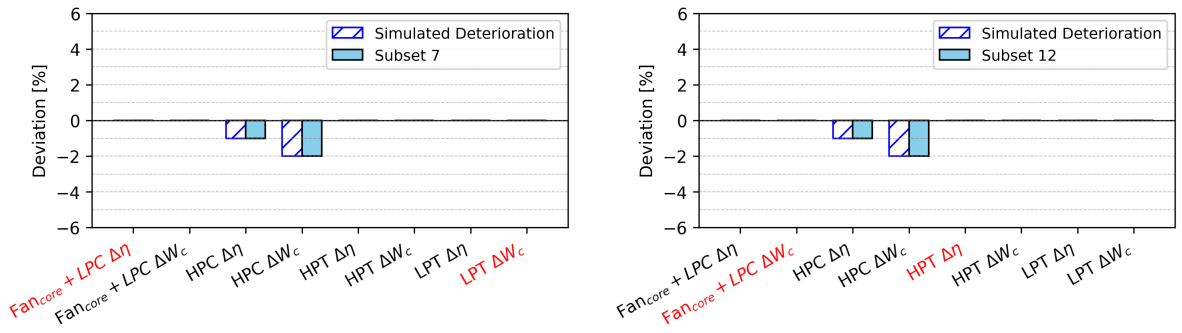

(a) Subset 7 **(b)** Subset 12

Figure 6.5: Subsets 7 and 12 with simulated HPC deterioration

Subsets that exclude the HPC flow capacity (subsets 1–3) exhibit a redistribution of deterioration to the Fan core and LPC (smearing). Similarly, subsets that exclude the HPC efficiency (subsets 4–6) show smearing of the degradation primarily to the Fan core and LPC efficiency (Figure 6.6).

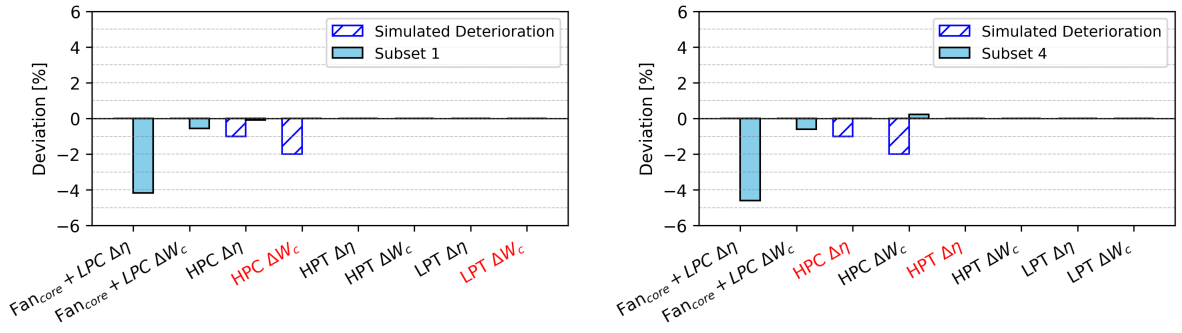

(a) Subset 1 **(b)** Subset 4

Figure 6.6: Subsets 1 and 4 with simulated HPC deterioration

6.2.3. HPT Simulated Deterioration

The HPT deterioration case follows trends similar to those observed for the compressor components. Subsets that include both HPT efficiency and flow capacity (subsets 1, 2, 5, 6, 7, and 8) accurately capture the imposed degradation. Two examples of these set are given in Figure 6.7a and Figure 6.7b.

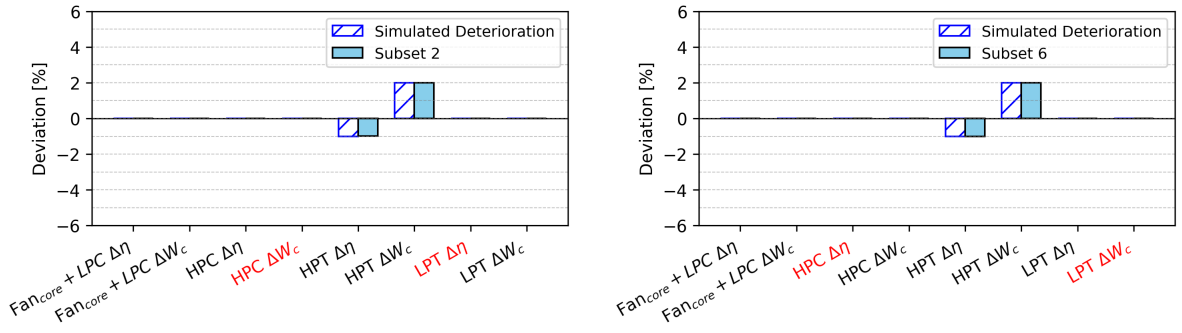

(a) Subset 2 **(b)** Subset 6

Figure 6.7: Subsets 2 and 6 with simulated HPT deterioration

When the HPT efficiency is excluded (subsets 3, 4, and 9), the model still predicts the flow capacity change with reasonable accuracy. However, the efficiency degradation is redistributed to the LPT,

affecting both its efficiency and flow capacity estimates (Figure 6.8). This again demonstrates the smearing effect resulting from incomplete representation of component health indicators.

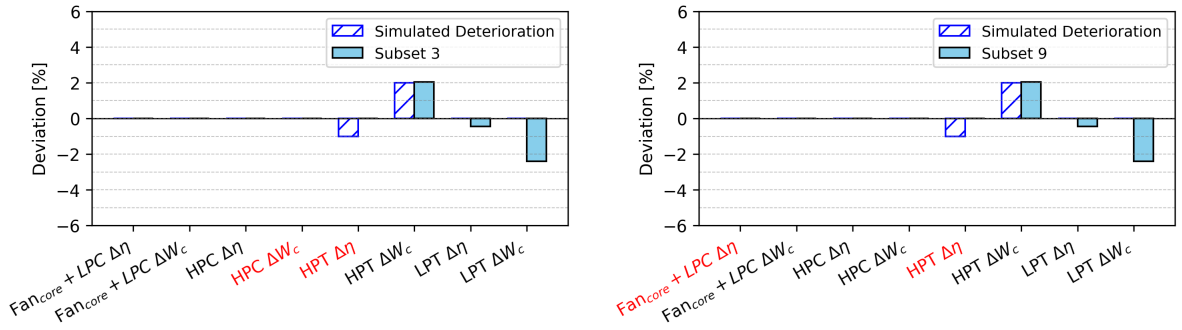

(a) Subset 3 **(b)** Subset 9

Figure 6.8: Subsets 3 and 9 with simulated HPT deterioration

6.2.4. LPT Simulated Deterioration

The LPT deterioration case exhibits three distinct trends. Subsets that include both LPT efficiency and flow capacity (subsets 3, 4, and 9) provide accurate predictions of the imposed degradation. Subsets that include only the LPT efficiency (e.g., subsets 1, 6, and 7) show a redistribution of degradation effects to the HPT, indicating strong coupling between turbine components, as can be seen in the example of Figure 6.9a. In addition, subsets 2, 5 and 6 only contain the LPT flow capacity health indicator, and this lead to a big jump in LPT flow capacity, shown in Figure 6.9b.

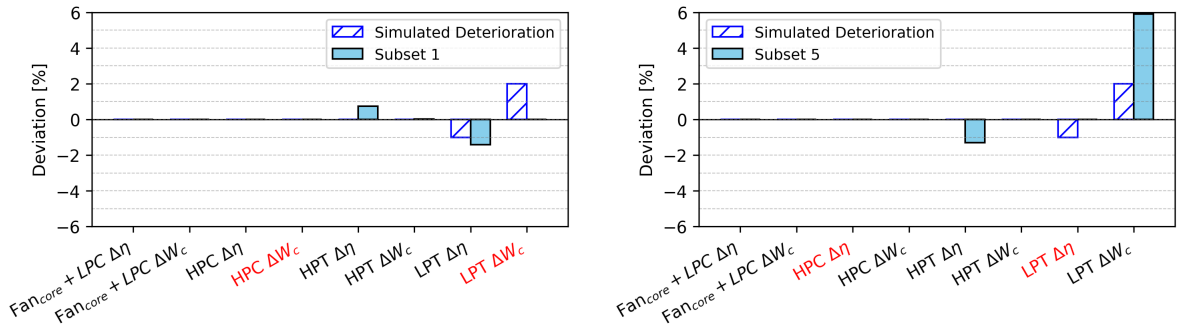

(a) Subset 1 **(b)** Subset 5

Figure 6.9: Subsets 1 and 5 with simulated LPT deterioration

6.2.5. Core Simulated Deterioration

Core deterioration, defined as simultaneous degradation of the HPC and HPT, is accurately predicted only by subsets that include all corresponding health indicators. In this study, subsets 7 and 8 satisfy this condition and can perfectly predict this core deterioration.

Other subsets fail to capture the core deterioration accurately, instead smearing the deterioration to the Fan core and LPC efficiency and the LPT flow capacity, as can be seen in subset 1 and 12, which are presented as an example in Figure 6.10 below.

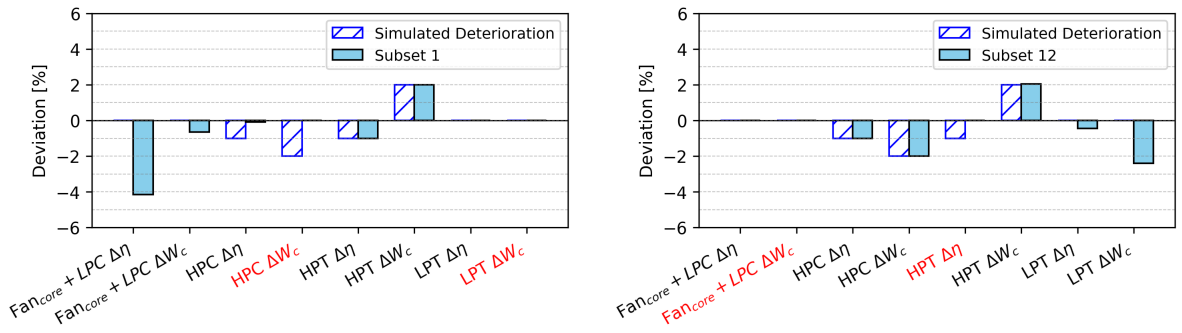

(a) Subset 1 **(b)** Subset 12

Figure 6.10: Subsets 1 and 12 with simulated core deterioration

6.2.6. Full Simulated Deterioration

Finally, full simulated deterioration was used to verify the subsets under full engine deterioration. However, since all health indicators deviate from their initial healthy state, and only 6 out of 8 health indicators can be used for each subset, there will always be smearing. Nevertheless, it is important to see if some subsets are able to correctly represent health indicators.

Despite this limitation, certain subsets perform better than others. The best-performing subsets exclude the Fan core and LPC efficiency, with the resulting smearing partially compensated by the HPT efficiency or the LPT flow capacity (Figure 6.11a and Figure 6.11b).

In contrast, the worst-performing subsets exclude the Fan core and LPC flow capacity, leading to significant and unrealistic redistribution of degradation to other components, particularly the HPC flow capacity.

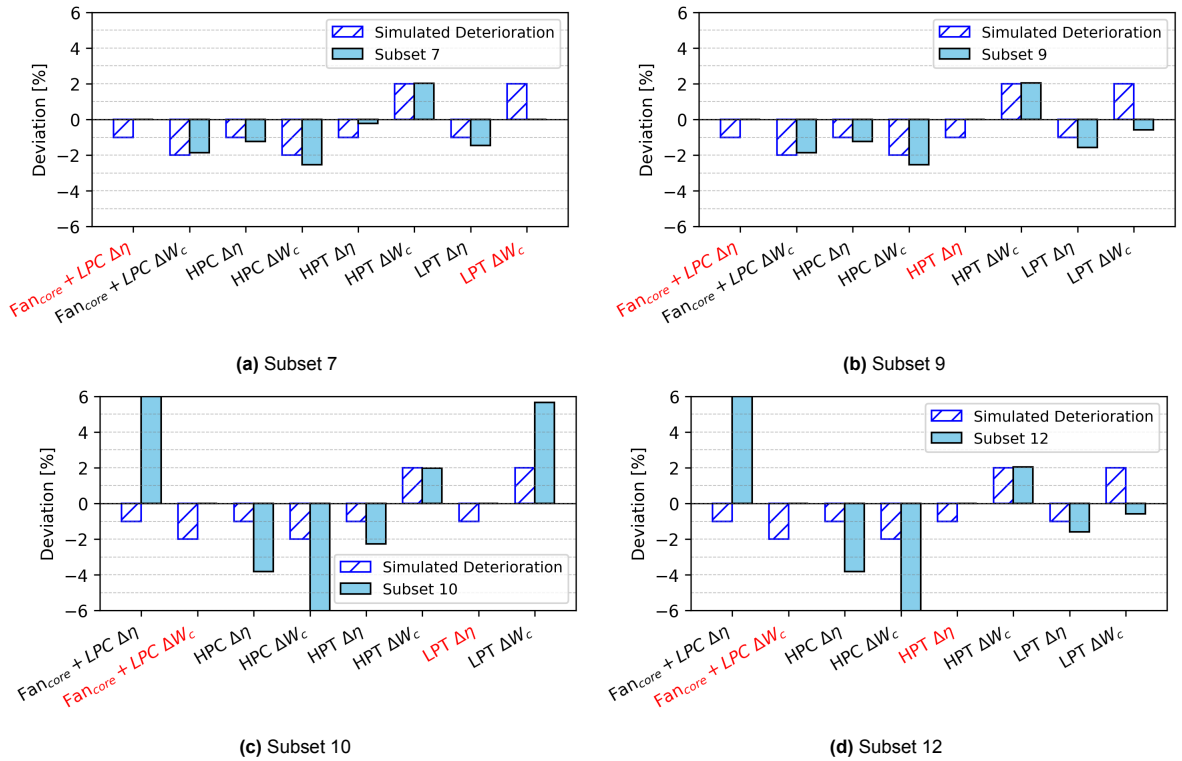

(a) Subset 7 **(b)** Subset 9

(c) Subset 10 **(d)** Subset 12

Figure 6.11: Subsets 7, 9, 10 and 12 with full simulated deterioration

6.3. Conclusion on Health Indicator Subsets

The analysis of the individual subsets demonstrates that no single subset is capable of accurately predicting all types of simulated deterioration. Accurate estimation of component degradation is only achieved when both corresponding health indicators are included in the subset.

Furthermore, each subset exhibits specific strengths and weaknesses in predicting certain deterioration cases. Therefore, a method has to be constructed to combine information from multiple subsets. This approach, referred to as the Component Exclusion Method (CEM), is introduced and discussed in the following chapter.

7

Component Exclusion Method (CEM)

This chapter presents the Component Exclusion Method. First, its concept is presented, whereafter the weight-factor matrix is explained, as well as the health indicator bounds. Then, the training and testing procedures of the CEM are presented, and three possible extensions to the Component Exclusion Method are displayed.

7.1. Concept

The Component Exclusion Method (CEM) is developed to combine the outputs of multiple subset health indicator prediction into a single, robust estimate of the engine health indicators. As previously discussed, individual subsets exhibit varying levels of accuracy for different health indicators. The central idea of the CEM is therefore to exploit this behavior by assigning appropriate weights to each subset, such that the combined estimate is optimized.

The methodology consists of the following steps:

1. Identification of subsets that reliably converge when applied to on-wing data.
2. Construction of a representative dataset of simulated deterioration cases, divided into training and test sets, reflecting realistic on-wing degradation patterns.
3. Adaptive Modeling simulations using the various subsets found in step 1.
4. Optimization of a weight-factor matrix using the simulated deterioration training dataset.
5. Verification of the weight-factor matrix using the simulated deterioration test dataset.
6. Execution of adaptive modeling simulations for each selected subset on on-wing data.
7. Combination of subset predictions using the weight-factor matrix to obtain final health indicator estimates.

The framework of the CEM is best visualized using a flowchart, Figure 7.1. Since test sets are introduced solely for the purpose of evaluating the performance of the method, they do not form an integral part of the operational workflow of the CEM itself. In practical applications, the methodology relies exclusively on the training phase for determining the weight-factor matrix, after which the framework can be directly applied to on-wing data without requiring further validation steps.

It is important to emphasize that the training and testing cases are essential during the development phase, as it ensures that the optimized weight-factor matrix does not overfit to specific simulated deterioration cases. By assessing its performance on independent test cases, the CEM can prove capable of generalizing to previously unseen degradation conditions and, consequently, to real engine behavior. Nevertheless, once this has been demonstrated, the test dataset no longer contributes to the estimation procedure and can therefore be excluded from the final implementation framework.

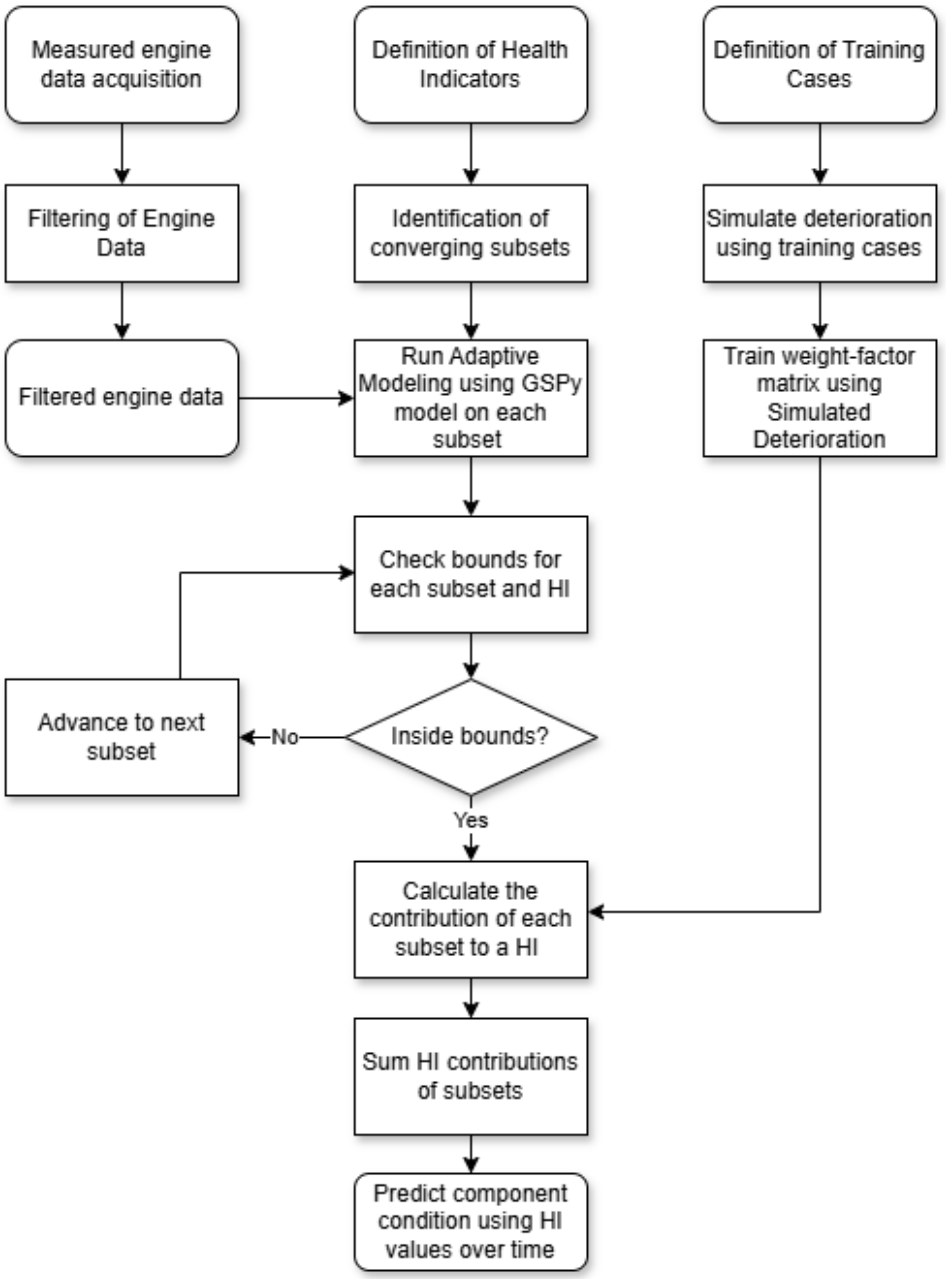

Figure 7.1: CEM flowchart

7.2. Weight-Factors

A key element of the CEM is the weight-factor matrix, which defines how the outputs of individual subsets are combined. Each subset s provides a prediction for simulated deterioration case c and health indicator h:

$$X_{s,c,h} \tag{7.1}$$

For a fixed health indicator h and simulated deterioration case c, the subset predictions can be written as a vector containing all subset predictions:

$$\mathbf{X}_{c,h} = \begin{bmatrix} X_{1,c,h} \\ X_{2,c,h} \\ \vdots \\ X_{S,c,h} \end{bmatrix} \tag{7.2}$$

where S denotes the total number of subsets. The weight-factor matrix assigns a weight to each subset for each health indicator, thereby enabling a weighted-combination of these predictions. The matrix is then defined as

$$\mathbf{W} = \begin{bmatrix} w_{1,1} & w_{1,2} & \dots & w_{1,H} \\ w_{2,1} & w_{2,2} & \dots & w_{2,H} \\ \vdots & \vdots & \ddots & \vdots \\ w_{S,1} & w_{S,2} & \dots & w_{S,H} \end{bmatrix} \tag{7.3}$$

where:

- S is the number of subsets,
- H is the number of health indicators,
- $w_{s,h}$ represents the importance of subset s for health indicator h.

Each column of $\mathbf{W}$ corresponds to one health indicator. For each simulated deterioration case c and health indicator h, the CEM estimate is computed as a weighted sum:

$$\hat{X}_{c,h} = \sum_{s=1}^{S} w_{s,h}\, X_{s,c,h} \tag{7.4}$$

Defining the weight vector for health indicator h as

$$\mathbf{w}_h = \begin{bmatrix} w_{1,h} \\ w_{2,h} \\ \vdots \\ w_{S,h} \end{bmatrix} \tag{7.5}$$

The weight matrix is obtained by minimizing the squared estimation error over all simulated deterioration cases and health indicators:

$$\min_{\mathbf{W}} \sum_{c=1}^{C} \sum_{h=1}^{H} \left(X_{c,h}^{\text{true}} - \hat{X}_{c,h} \right)^2 \tag{7.6}$$

where C denotes the number of degradation modes. Large values of $w_{s,h}$ increase the influence of subset s for health indicator h, whereas small values suppress its contribution. In addition, the minimize function contains a spread penalty, which increases if the spread of weight is very low. This ensures not all weight is given to one subset, but that preferable weights are spread over multiple subsets.

7.3. Health Indicator Bounds

To ensure physically meaningful results, bounds are introduced to filter out unrealistic subset predictions. In some cases, individual subsets produce excessively large deviations from the nominal condition for specific health indicators. Such predictions are considered non-physical and may arise from poor observability, numerical instability, or due to high measurement uncertainty.

During training, a set of predefined bounds is imposed on the predicted health indicators to prevent these outliers from influencing the training process. The selected bounds, shown in Figure 7.2, reflect realistic limits on component deterioration based on expected health indicator behavior.

Mathematically, for each subset s, simulated deterioration case c, and health indicator h, the predicted value $X_{s,c,h}$ is only retained if it satisfies

$$X_h^{min} \leq X_{s,c,h} \leq X_h^{max}, \tag{7.7}$$

where X_h^{min} and X_h^{max} denote the lower and upper bounds for health indicator h. Predictions that violate these bounds are excluded from the optimization problem used to determine the weight-factor matrix using the following function:

$$\delta_{s,c,h} = \mathbf{1}\left\{X_h^{\min} \leq X_{s,c,h} \leq X_h^{\max}\right\} \tag{7.8}$$

This function is then used to modify the objective function of the weight-factor matrix optimization as:

$$\min_{\mathbf{W}} \sum_{c=1}^{C} \sum_{h=1}^{H} \delta_{s,c,h} \left(X_{c,h}^{\text{true}} - \hat{X}_{c,h}\right)^2 \tag{7.9}$$

Where $\delta_{s,c,h}$ ensures only valid subset contributions are considered in the error minimization.

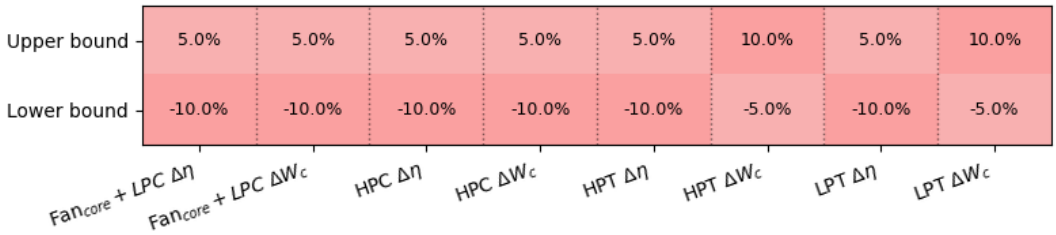

Figure 7.2: Deterioration bounds set

7.4. CEM Training

The weight-factor matrix is obtained through a supervised training procedure using simulated deterioration data. The training dataset is designed to reflect realistic on-wing deterioration behavior, with particular emphasis on degradation modes that are frequently observed in practice, such as HPT deterioration. This led to the training cases seen in Figure 7.3.

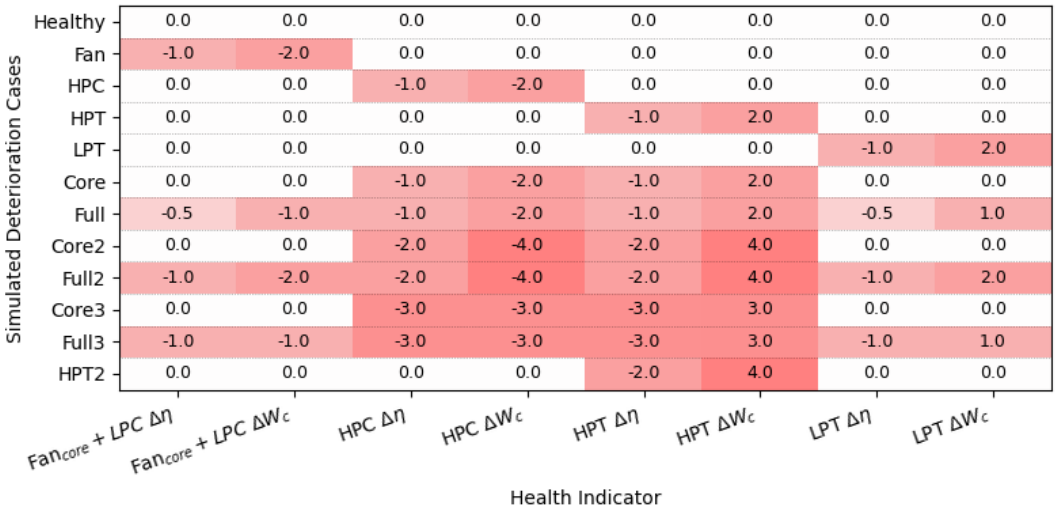

Figure 7.3: Simulated deterioration training cases: Each value representing the percentage change of a certain health indicator

Off-design simulations are ran using these simulated deterioration, resulting in gas path parameters deviating from a nominal off-design simulation. These outputs are then used as inputs in the Adaptive Modeling simulation done using health indicators found in each subset.

7.5. CEM Testing

Following training, the performance of the CEM is evaluated using an independent test dataset. This step is essential to verify that the weight-factor matrix does not overfit specific deterioration patterns. The test cases differ from the training cases but maintain similar characteristics in terms of dominant deteriorated health indicators, particularly core deterioration. This ensures that the evaluation remains representative of realistic engine behavior while still providing a meaningful assessment of the predictive accuracy. The simulated test cases are presented in Figure 7.4, where each case represents a distinct combination of health indicator deviations.

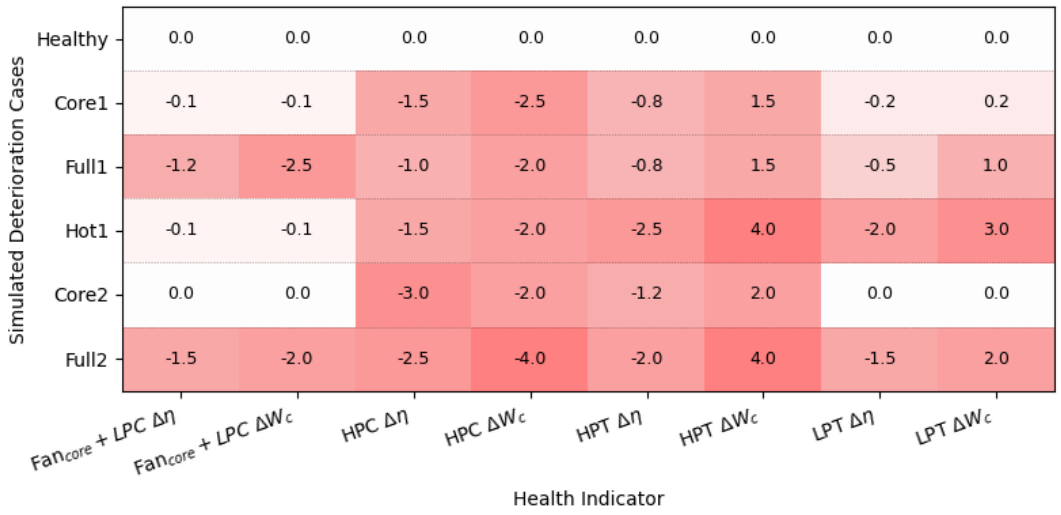

Figure 7.4: Simulated deterioration test cases: Each value representing the percentage change of a certain health indicator

The training and testing procedure are done for three operating points: Takeoff, climb and cruise. The operating conditions for these points are the average N1k, Mach number and altitude for these flight phases over a large set of aircraft, as illustrated in Figure 9.1. This ensures that the weights are trained for a specified operating range, as subsets might behave differently at takeoff than at cruise.

7.6. Extensions to the Component Exclusion Method

This section presents several extensions to the Component Exclusion Method (CEM), introduced to address limitations identified during preliminary application. These extensions aim to improve the robustness and accuracy of the health indicator estimation. The impact of these modifications is evaluated in the verification chapter.

7.6.1. Handling of LPT Flow Capacity Inaccuracy

During preliminary application of the Component Exclusion Method, it was observed that certain health indicators could not be reliably estimated. In particular, the LPT flow capacity consistently showed poor predictive accuracy, remaining close to zero even in the presence of simulated deterioration. This behavior can be attributed to limited observability of the low-pressure turbine. Within the available sensor set, only the inlet temperature of the LPT is measured, providing insufficient information to accurately estimate deterioration in the LPT flow capacity.

To improve the robustness of the method, the LPT flow capacity is excluded from the set of health indicators. This reduces the total number of health indicators from eight to seven. Consequently, the CEM is reformulated such that six out of the seven remaining health indicators are estimated in each subset, resulting in seven possible subsets.

Figure 6.1 illustrates the subsets that exclude the LPT flow capacity. Among these, only subsets 1, 7, and 11 consistently converge for both simulated deterioration cases and on-wing data. These subsets are therefore selected for further analysis to evaluate the impact of excluding the LPT flow capacity on the estimation accuracy.

7.6.2. Combining Low Pressure Efficiencies

In addition to the LPT flow capacity, the efficiencies associated with the low-pressure shaft, specifically the fan core + LPC efficiency and the LPT efficiency, were found to exhibit large variability across different subsets. This indicates that these parameters are not independently observable given the available measurements.

This limitation is primarily related to the sensor configuration on the low-pressure shaft, where the available measurements provide insufficient information to accurately estimate these efficiencies. As a result, the estimation of these parameters becomes highly sensitive and leads to non-physical results.

To address this issue, a coupling strategy between the low-pressure efficiencies is introduced. Rather than defining a single combined health parameter, the method consists of modifying the subsets and adding a residual equation. All subsets are constructed such that both the fan core + LPC efficiency and the LPT efficiency are always included. Each subset is then completed by adding one additional health indicator, resulting in a total of six subsets. An additional residual equation is also created, which equals the difference between the fan core + LPC efficiency and the LPT efficiency. By doing it in this way, the solver tries to set the efficiencies equal but has some space to adapt them.

Both these extensions will be implemented and tested. Consequently, its results will be compared to the CEM without any of these extension implemented to see if these extensions can benefit the prediction accuracy, robustness or computational efficiency.

Part III

Verification & Validation

8

Verification with Simulated Deterioration

This chapter evaluates the performance of the Component Exclusion Method (CEM) using simulated deterioration cases. The primary objective is to assess both the prediction accuracy and robustness of the method. Prior to this evaluation, a sensitivity analysis is conducted to quantify the influence of deterioration in health indicators on gas path parameters. In addition, before performing verification using simulated deterioration, it is necessary to identify which subsets can be reliably applied to on-wing data. As discussed in the previous chapter, several subsets were found to be ill-conditioned, leaving twelve subsets that consistently converge under simulated deterioration conditions. However, not all of these subsets demonstrate stable convergence when applied to on-wing data. Therefore, an initial analysis is performed to assess their convergence in on-wing operating conditions.

8.1. Sensitivity Analysis

The sensitivity analysis was performed to verify the effect of deterioration in components on the gas path parameters, providing insight into the sensitivity of gas path parameters to changes in health indicators. Each health indicators is individually subjected to a -1% deterioration. For the turbine flow capacity, this corresponds to a +1% increase in flow capacity due to the inverse relationship between flow capacity and deterioration. Using the simulation deterioration class implemented in GSPy, as described in subsection 5.2.1, off-design simulations are performed. The results are then compared against a *healthy* GEnx-1B simulation. This investigation does assume a linear relationship, which might not be true with higher deterioration cases.

An illustrative example is presented in Figure 8.1. The schematic of the GEnx-1B engine model is shown alongside the corresponding sensor locations. The inlet conditions, defined by PT2, TT2, and N1, are kept constant. A −1% deterioration in high-pressure turbine (HPT) efficiency is introduced, resulting in changes in gas path parameters relative to the healthy condition.

Specifically, the turbine outlet temperature (TT49) increases as a consequence of reduced HPT efficiency, as less work is extracted for a given enthalpy drop. To maintain a constant spool speed (N1), the engine must compensate for this loss in efficiency. This is achieved by increasing the fuel flow, thereby raising the overall enthalpy level in the cycle. As a result, both the energy available to the low-pressure turbine (LPT) and the required work output are maintained

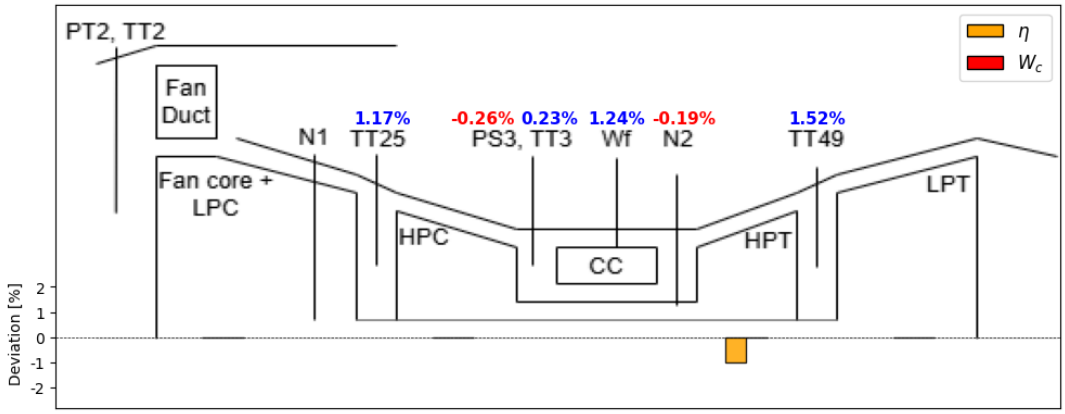

Figure 8.1: Gas path parameter changes due to HPT efficiency deterioration

The results of all sensitivity simulations are presented in Figure 8.2. These results provide valuable insight into the response of the gas path parameters to component deterioration. For example, most gas path parameters exhibit only minor changes in response to deterioration in the combined fan core and LPC efficiency. This indicates a low sensitivity to this health indicator, which in turn suggests limited predictive capability. This behavior also explains the frequent occurrence of smearing associated with this parameter. A similar trend is observed for the HPC flow capacity, where primarily N2 is affected while other gas path parameters remain relatively unchanged.

Furthermore, a decrease in component efficiency consistently results in an increase in fuel flow. This is expected, as additional fuel is required to compensate for the loss in efficiency and to maintain the required work output by increasing the available enthalpy in the cycle.

An interesting observation is the variation in TT25 resulting from LPT efficiency deterioration. Since N1 is maintained constant, one might expect TT25 to remain unchanged. However, this is not the case. A reduction in LPT efficiency implies that less work is extracted from the flow to drive the low-pressure spool. To maintain the required N1, the engine compensates by increasing fuel flow, which raises the turbine inlet temperature and leads to an increase in N2. As a result, the high-pressure

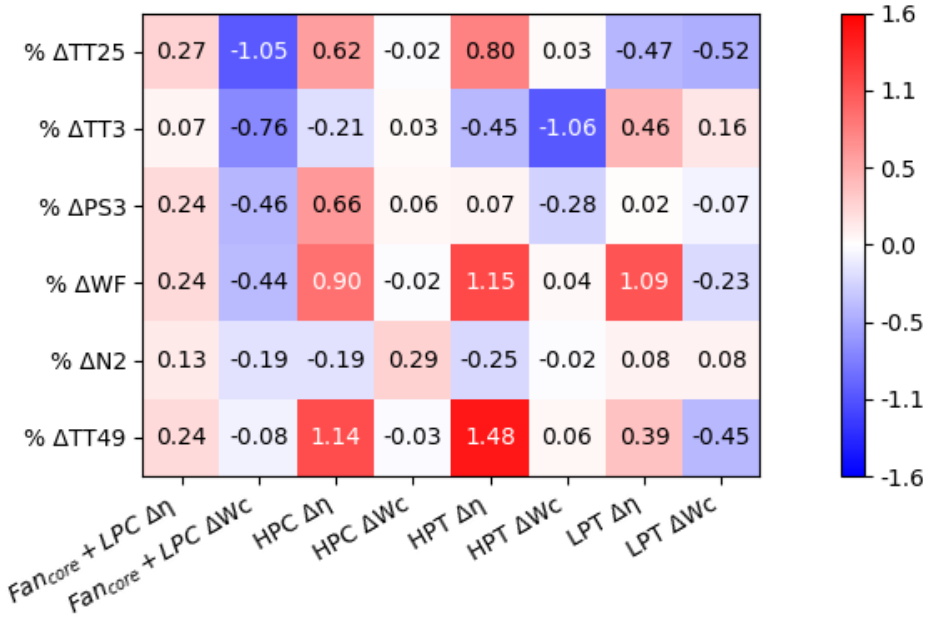

Figure 8.2: Sensitivity Analysis of gas path parameters; Change in gas path parameters is represented due to a 1% change in health indicators

compressor (HPC) operates at a higher work level, altering its corrected mass flow. To maintain flow continuity, the corrected mass flow at the outlet of the fan core and LPC must match that of the HPC inlet. Consequently, the operating point of the fan core and LPC shifts on its performance map, leading to a change in map efficiency and therefore TT25.

8.2. Noise

Sensor measurements are influenced by noise. As shown before in Table 2.3, the recommended accuracy of each sensor also differs. This can lead to inaccuracies in health indicator predictions. An investigation is therefore performed to find out how much this noise influences the outcome of CEM, and how much the implementation of measurement tolerances helps with this.

8.2.1. Methodology

Each sensor has its own noise, and this noise is unknown. However, the recommended system accuracies can be used to estimate this noise. It is assumed the system accuracies at twice the standard deviation (2σ) [24]. This means that approximately 95.5% of sensor measurements lie within that deviation.

This investigation aims to find out what the mean and standard deviations are of the health indicators predictions with noise. This is done using a Monte Carlo simulation, in which synthetic noise is added based on Gaussian distributions. For each of the 500 points, noise is added and on this data the CEM is performed. This amount of points is deemed enough as the mean and standard deviation have converged [24]. By performing a large number of such realizations, a statistical distribution of the predicted health indicators is obtained, allowing for the estimation of their mean, variance, and overall sensitivity to measurement uncertainty.

8.2.2. Results

The results indicate that incorporating measurement tolerances does not lead to an improvement in the health indicator distribution. A comparison between simulations performed with and without measurement tolerances shows no significant differences in the resulting health indicator distributions. All results presented correspond to simulations that include measurement tolerances, since it does slightly improve the convergence rate of on-wing data. Figure 8.3 illustrates the distributions obtained from the Monte Carlo simulations with measurement noise applied. It should be noted that the axis bounds differ between the fan core + LPC efficiency and LPT flow capacity plots, which should be considered when interpreting the spread of the results.

The mean values of most health indicators remain close to zero, as expected, since the simulated dataset does not include any component deterioration. An exception is observed for the fan core + LPC flow capacity, which shows a slight positive bias of approximately 0.45. This deviation suggests a systematic offset in the estimation of this parameter rather than an effect of measurement noise. This limitation can be addressed using baselining methods.

Although the true health state is nominal, the standard deviation of the estimated health indicators is non-zero due to the presence of measurement noise and model sensitivities. For the fan core + LPC efficiency and LPT flow capacity, the standard deviation is very low. This is a direct consequence of the CEM weight-factors (see section 8.4), which effectively forces these parameters to remain close to zero. For the remaining health indicators, the standard deviation lies between 1 and 2, indicating a noticeable spread in the estimates. This suggests that these parameters are more sensitive to measurement noise and less constrained by the model.

Overall, the results demonstrate that while measurement noise introduces variability in the estimated health indicators, the inclusion of measurement tolerances does not significantly mitigate this effect. Consequently, when analyzing health indicators over time, isolated deviations should be interpreted with caution. A single deviation is likely due to measurement noise rather than actual component deterioration, and therefore should not be considered as a sign of degradation. Instead, consistent trends or a shifts of a health indicator over multiple data points is required to confidently identify true deterioration.

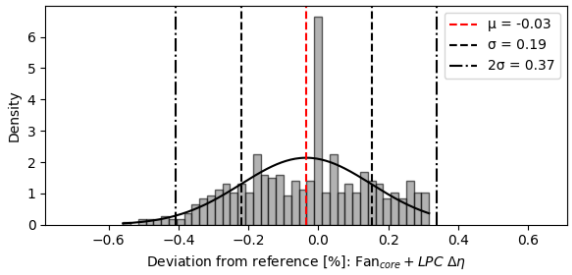

(a) Fan core + LPC efficiency distribution

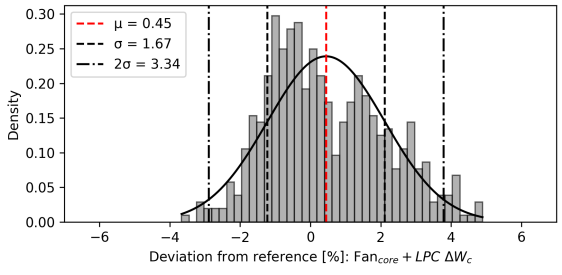

(b) Fan core + LPC flow capacity distribution

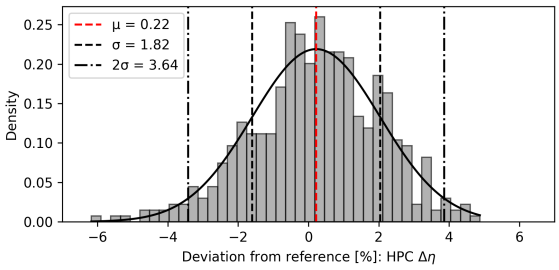

(c) HPC efficiency distribution

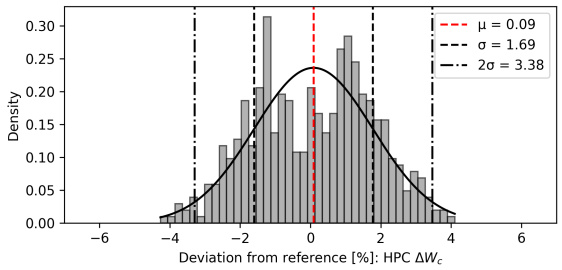

(d) HPC flow capacity distribution

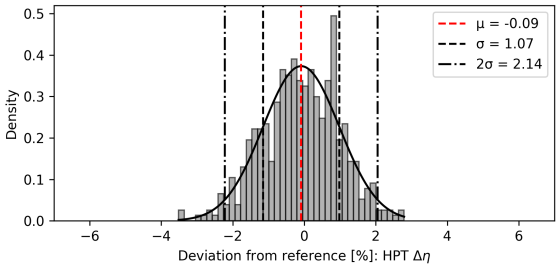

(e) HPT efficiency distribution

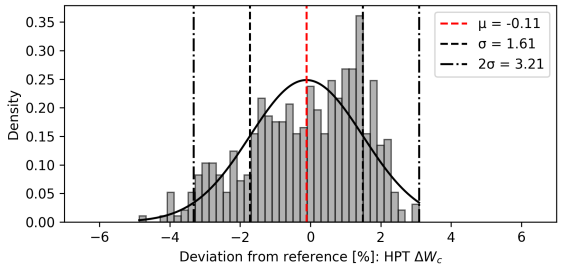

(f) HPT flow capacity distribution

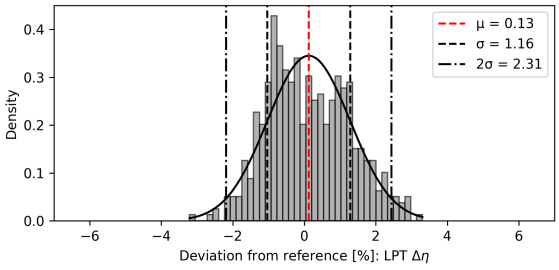

(g) LPT efficiency distribution

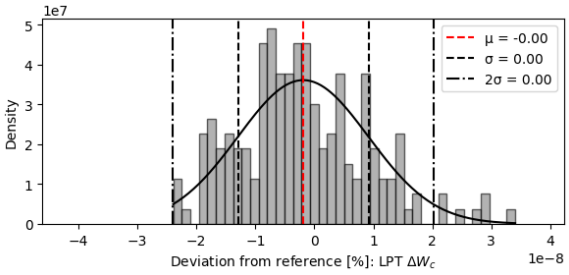

(h) LPT flow capacity distribution

Figure 8.3: Probability density distributions for each health indicator using a Monte Carlo simulation

8.3. Testing Convergence with On-Wing Data

To assess the applicability of CEM under realistic on-wing conditions, a sample of on-wing data points was analyzed. The convergence behavior of all twelve subsets was evaluated by performing adaptive modeling simulations. It was found that only six out of twelve possible subsets consistently converged. The convergence percentage of all twelve subsets is shown in Figure 8.4.

In addition to convergence reliability, computational performance was also considered. Non-converging cases were found to require significantly longer computation times compared to converged cases, further reducing their practical applicability.

Based on these observations, only the six consistently converging subsets are retained for further analysis. These subsets define the baseline configuration of the CEM. Consequently, the weight-factor matrix is trained and evaluated using only these subsets.

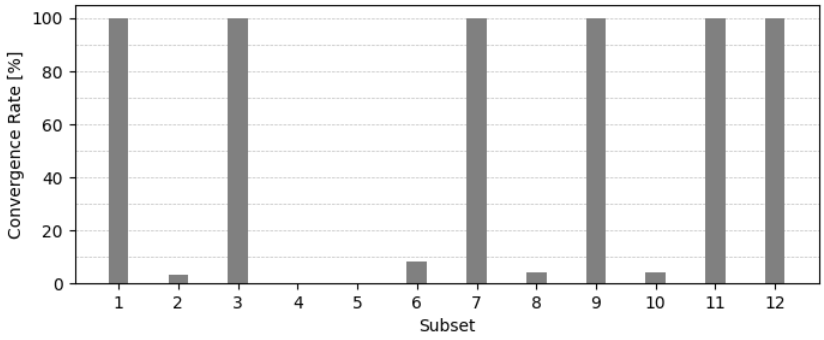

Figure 8.4: Adaptive Modeling Convergence Percentage of an On-Wing Data Sample

8.4. Weight-Factor Matrix Training

The first step consists of training the weight-factor matrix using the simulated deterioration cases and the six subsets defined the previous section. An illustrative example of the resulting weight-factor matrix is presented in Figure 8.5.

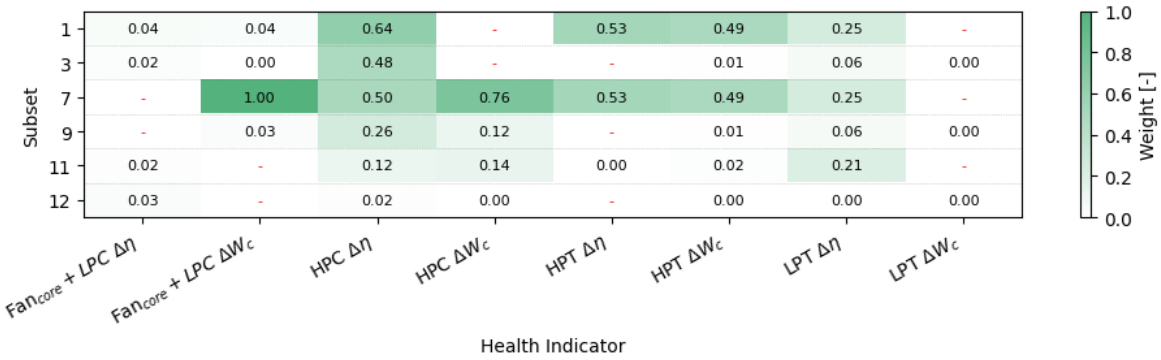

Figure 8.5: Weight-Factor Matrix; Values in cell represent the weight given to that subsets for a certain HI

It is important to emphasize that the values in the weight-factor matrix should not be interpreted as direct measures of the prediction accuracy of a given subset for a specific health indicator. Due to the elimination of subsets outside the bounds for certain cases, this can lead to a subset being excluded from the weight-factor matrix training for that case. For example, all subsets can accurately predict the HPT flow capacity. However, subsets 3, 9, 11 and 12 are excluded for some cases of full engine deterioration due to one or more health indicators exceeding bounds. If all subsets are given equal weight, this will lead to a lower prediction value, since the influence of excluded subsets is not accounted for in the final weighting.

The trained weight-factor matrix highlights a couple important results. The fan core + LPC efficiency cannot be predicted accurately, and therefore the combined weight of all subsets is only 0.12. This will lead to a low prediction of fan core + LPC efficiency, even if deterioration is present. This is consistent with the results of the sensitivity analysis, as deterioration in the fan core + LPC does not lead to big changes in gas path parameters.

In addition, subset 7 is given almost the full weight of predicting the fan core + LPC flow capacity. Although spread of weights is favorable and is incorporated in the weight-factor matrix optimization, this weight distribution led to the lowest RMSE.

The HPC efficiency weights are spread the most. With all subsets having a non-zero weight. Furthermore, the HPC flow capacity can be effectively combined using subsets 7, 9 and 11. The HPT efficiency is only included in subsets 1, 7 and 11. Due to subset 11 exceeding bounds for a number of simulated deterioration cases, subsets 1 and 7 are given the weight.

Finally, the LPT efficiency can be accurately predicted by most health indicators, while the LPT flow capacity is given zero weight. This means that none of the subsets can give predictions matching the simulated deteriorated cases, and the LPT flow capacity will remain at zero.

8.5. Results

This section presents the results obtained from the training and testing of the Component Exclusion Method (CEM) configuration described in the previous chapter.

8.5.1. Component-Level Training Cases

As illustrated in Figure 7.3, the performance of the CEM is evaluated by creating separate simulated deterioration cases. The initial four cases represent component-level deterioration cases. The corresponding results after implementing the CEM are presented in Figure 8.6.

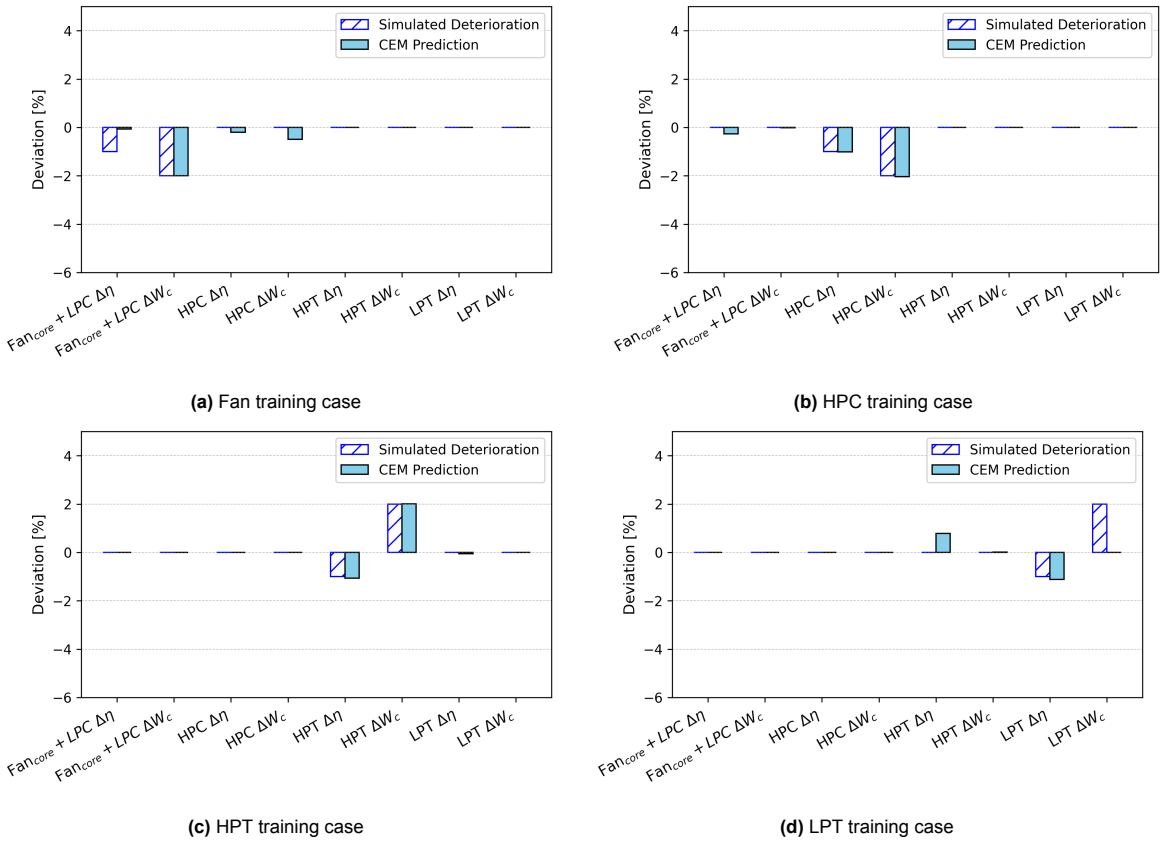

(a) Fan training case

(b) HPC training case

(c) HPT training case

(d) LPT training case

Figure 8.6: Simulated deterioration prediction of CEM using component-level deterioration cases

Overall, the CEM demonstrates strong capability in predicting and quantifying component-level deterioration. In particularly, deterioration of the core components can be predicted accurately. However, systematic prediction error can be observed for some health parameters. The fan core + LPC efficiency is underestimated, which is in line with the expectations based on the weight-factor matrix. A similar effect is observed for the LPT flow capacity. This behavior shows that the coupling of the low-pressure components cannot be fully solved by the current AM configuration.

8.5.2. Multi-Component Training Cases

In addition to component-level cases, the weight-factor matrix is trained using multi-component deterioration scenarios, including both core and full engine degradation. The results presented in Figure 8.7 demonstrate that the CEM is capable of accurately predicting the majority of health indicators under these more complex conditions.

Specifically, all health indicators are well predicting in magnitude, with the exception of the Fan core

and LPC efficiency and the LPT flow capacity. These two parameters remain challenging to estimate accurately, consistent with observations from the component-level cases. Despite these limitations, the CEM shows strong performance in predicting the distribution and magnitude of deterioration across the engine.

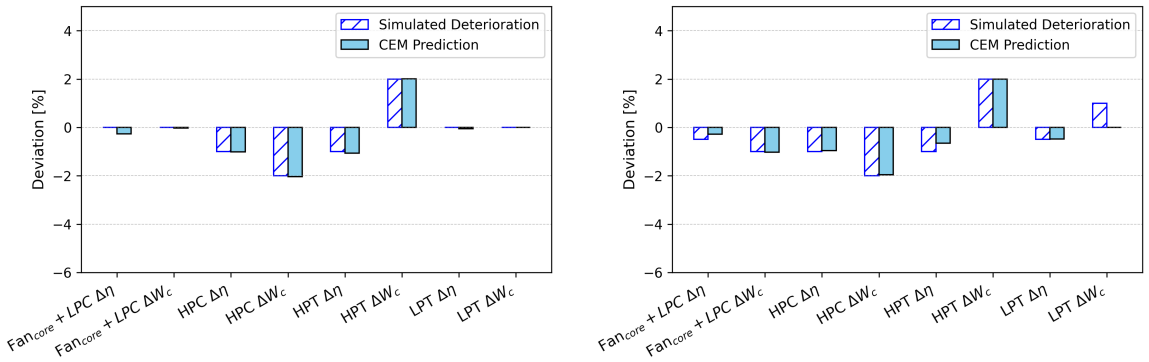

(a) High core deterioration training case

(b) High full deterioration training case

Figure 8.7: Simulated deterioration prediction of CEM using multi-component deterioration cases

8.5.3. Test Cases

The test cases are used to verify that the CEM is also capable of predicting deterioration cases that differ from the training cases. The results of the test cases are presented in Figure 8.8.

The first set of test cases, representing low core and low full engine deterioration are presented in Figure 8.8a and Figure 8.8b. These cases closely resemble the corresponding training cases. The results indicate that the CEM is not overfitted, as it is able to accurately detect and quantify variations in deterioration compared to the exact training conditions.

Figure 8.8c and Figure 8.8d highlight medium deterioration cases. These indicate that the CEM can correctly identify which components have deteriorated, although some discrepancies in magnitude estimation are visible. This is especially visible in the case of core deterioration, where the HPC efficiency deterioration is overestimated.

Finally, Figure 8.8e shows simultaneous deterioration of the HPC, HPT and LPT. This case confirms that the CEM can correctly identify the components that have deteriorated. Nevertheless, the LPT flow capacity deterioration is smeared to the HPT efficiency.

In summary, the CEM is capable of accurately predicting the majority of health indicators across a range of unseen deterioration scenarios. However, consistent inaccuracies remain for the Fan core and LPC efficiency and the LPT flow capacity. The following sections address potential strategies to mitigate these limitations.

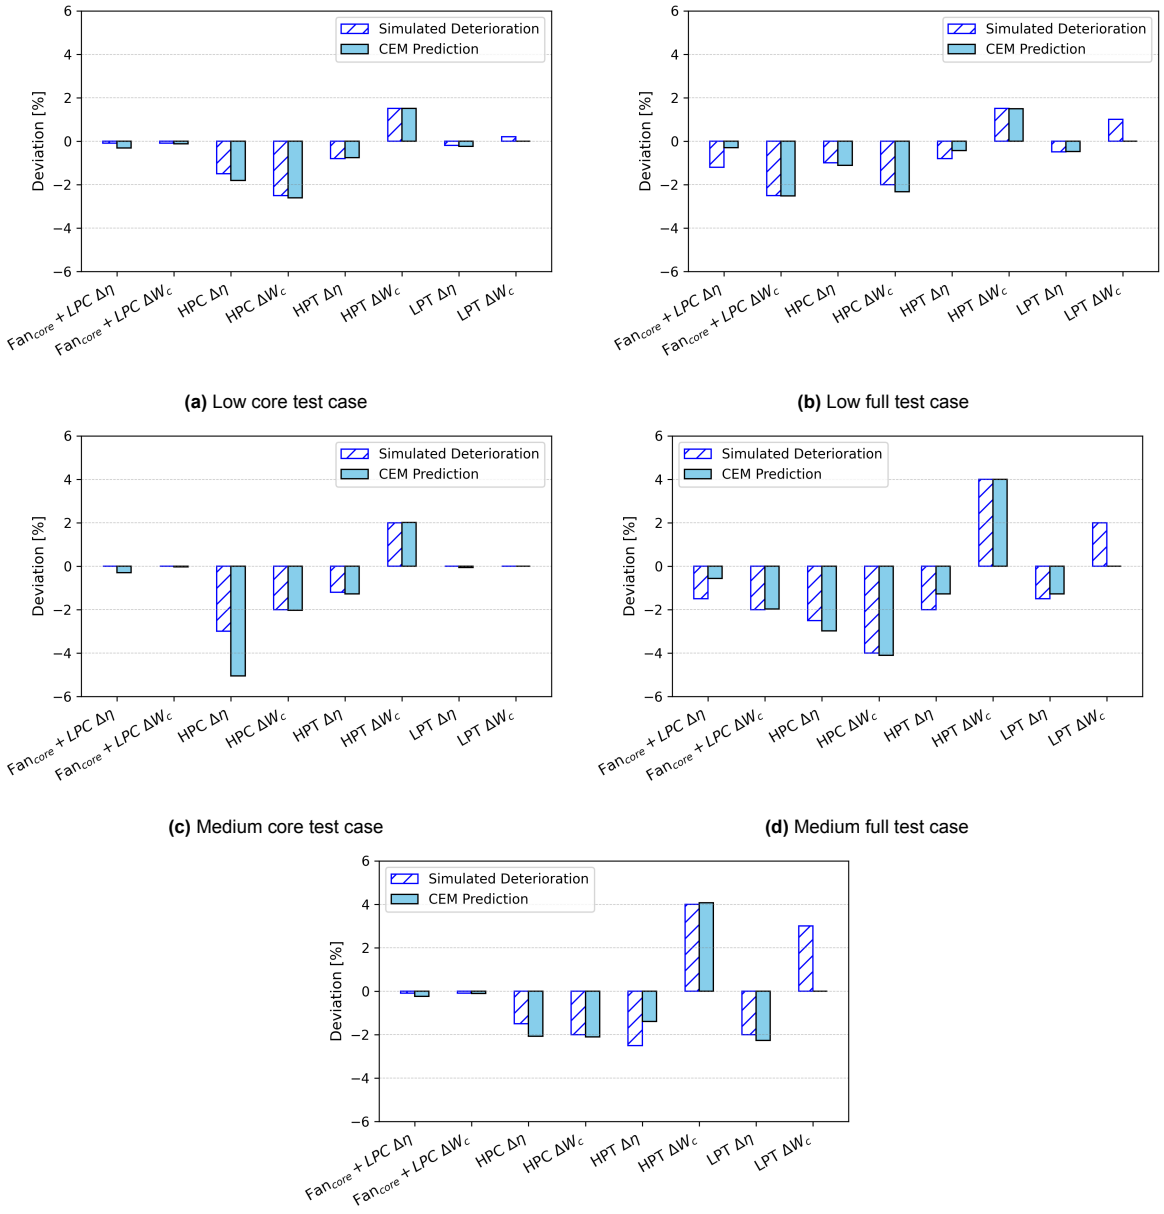

(a) Low core test case

(b) Low full test case

(c) Medium core test case

(d) Medium full test case

(e) HPT & LPT test case

Figure 8.8: Simulated deterioration prediction of CEM using test cases

8.6. Extensions to the Component Exclusion Method Results

As explained in the previous chapter, several extensions to the Component Exclusion Method were introduce to mitigate certain limitation, such as the LPT Flow Capacity and Fan core + LPC efficiency health indicator predictions. The results of those extensions are shown in the following section, followed by a conclusion into the highest performing CEM configuration.

8.6.1. Handling of LPT Flow Capacity Inaccuracy

The results obtained from the original CEM configuration with six subsets indicated that the LPT flow capacity is the least accurately estimated health indicator. Considering that the LPT typically exhibits limited deterioration under on-wing operating conditions [24], an investigation was performed where the LPT flow capacity was not taken into consideration.

From the available subset combinations, several subsets do not include the LPT flow capacity. How-

ever, only part of these subsets demonstrate convergence across simulated deterioration cases. Furthermore, one of these subsets fails to produce stable results when applied to on-wing data. Therefore, only subsets 1, 7 and 11 are selected for further analysis on the impact of excluding the LPT flow capacity.

The training and testing procedures remain equal with those described previously. However, the deterioration cases are adapted to match the assumption that no deterioration occurs in the LPT flow capacity. The resulting subset predictions demonstrate that there are still discrepancies in the subset predictions. Two test cases are presented in Figure 8.9.

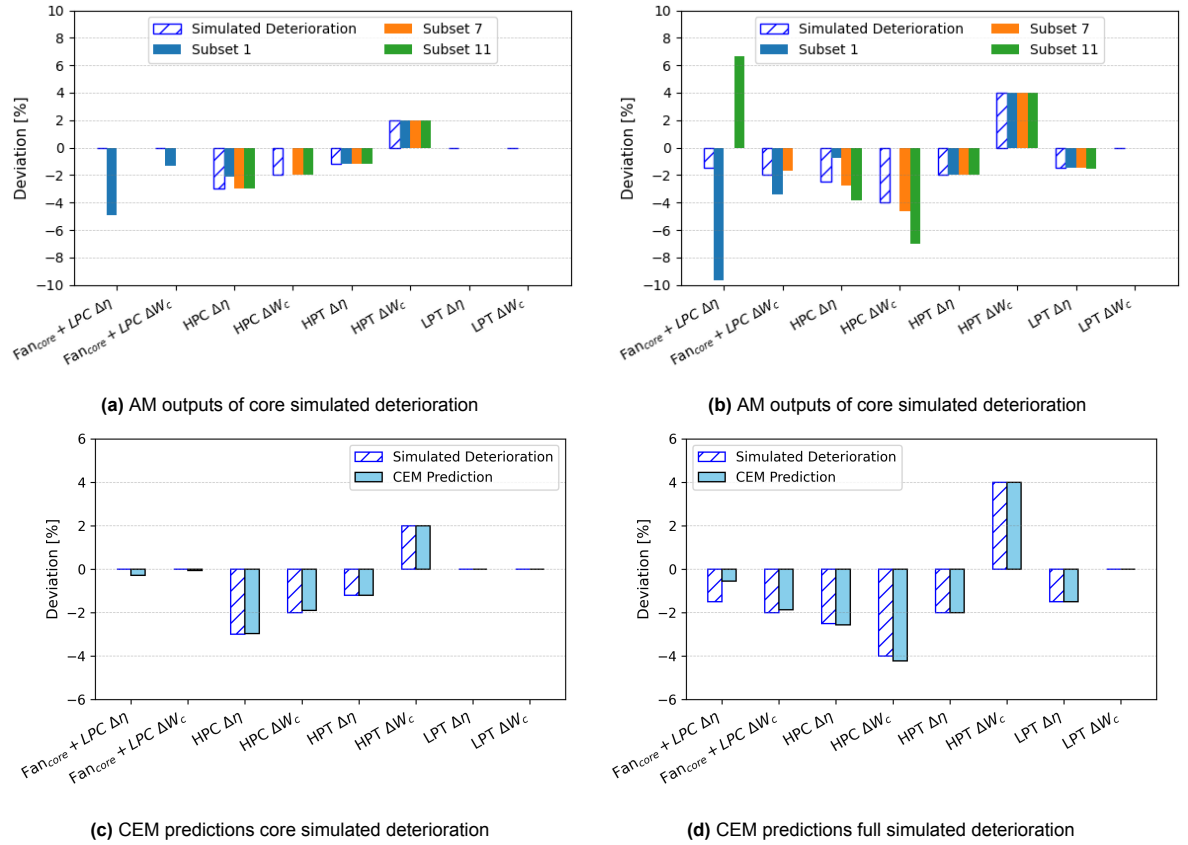

(a) AM outputs of core simulated deterioration

(b) AM outputs of core simulated deterioration

(c) CEM predictions core simulated deterioration

(d) CEM predictions full simulated deterioration

Figure 8.9: CEM results of converging subsets without the LPT flow capacity

The resulting root mean squared error (RMSE) values (Figure 8.10) demonstrate that the reduced three-subset configuration achieves comparable, if not better performance than the original six-subset configuration. A notable influence is the HPC efficiency, where the three-subset configuration has a substantially lower RMSE. This deviation is attributed to a specific case in which the deterioration of the HPC efficiency is -2%, and the CEM predicts -4%. Excluding this outlier, the overall performance of the two configurations is fairly equal.

Notably, the RMSE values for the HPT and LPT health indicators are close to zero, indicating that these parameters can be predicted with near-perfect accuracy when the LPT flow capacity is excluded. This may suggests that the inclusion of the LPT flow capacity introduces uncertainties in the estimation of the other health indicators.

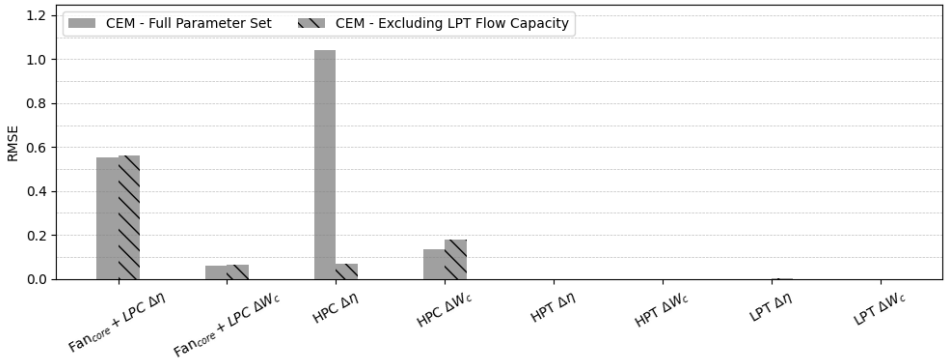

Figure 8.10: RMSE comparing the original CEM configuration with CEM - Excluding LPT Flow Capacity

In summary, the results indicate that, under the assumption of negligible LPT flow capacity degradation, the CEM can be simplified to a three-subset configuration without significant loss of predictive accuracy. Moreover, this configuration will lead to increased computational efficiency, as only three instead of six subsets will need to perform AM simulations.

8.6.2. Combining Low Pressure Efficiencies

An alternative refinement involves combining the efficiencies of low-pressure components within the CEM framework. This approach follows the same training and testing methodology as described previously, with specifically designed train and test sets to reflect the assumptions made. The performance of this approach, evaluated using simulated deterioration test cases, is presented in Figure 8.11.

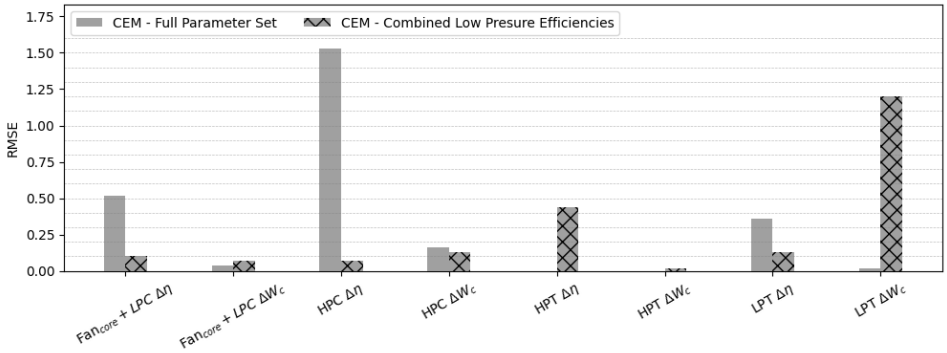

Figure 8.11: RMSE comparing the original CEM configuration with CEM - Combined Low Pressure Efficiencies

The findings indicate that combining low-pressure efficiencies yields mixed effects on prediction accuracy. Improvements are observed for the fan core + LPC efficiency, LPT efficiency, and HPC efficiency, suggesting that constraining the low-pressure component health efficiencies does remove some of the uncertainty present. The HPC efficiency high RMSE comes from the same discrepancy as before, where 1 test case leads to a HPC efficiency prediction of -4%, where the actual deterioration is -2%. Excluding this case gives similar HPC efficiency RMSE. However, this simplification comes at the cost of reduced accuracy for the HPT efficiency and the LPT flow capacity.

8.6.3. Combined Refinement Approach

Finally, both modifications, excluding the LPT flow capacity and combining low-pressure component efficiencies, were applied simultaneously. This resulted in a single subset configuration that excludes the LPT flow capacity while incorporating combined efficiency parameters. The RMSE results for this combined approach are shown in Figure 8.12.

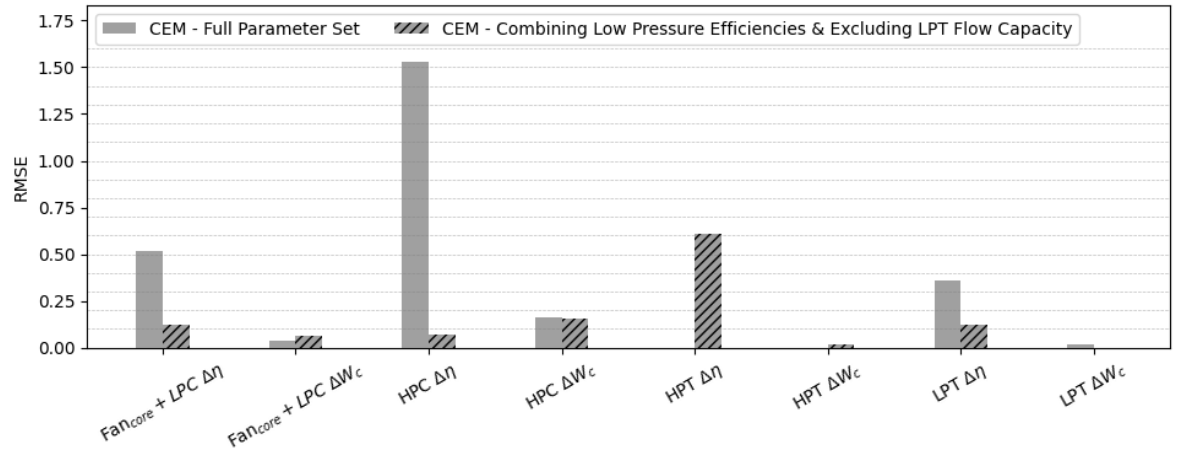

Figure 8.12: RMSE comparing the original CEM configuration with CEM - Combining Low Pressure Efficiencies & Excluding LPT Flow Capacity

The observed trends are consistent with those of the combined efficiency method. Most health indicators exhibit improved prediction accuracy, with the exception of the HPT efficiency, which shows a slight degradation in performance. This is however the most deteriorated component, thus high accuracy is essential.

In addition to predictive accuracy, the computational performance of each configuration is evaluated by measuring the runtime per data point using a QAR sample. The results, presented in Figure 8.13, indicate that the configurations excluding the LPT flow capacity improve the computational time, primarily due to the smaller number of subsets and reduced parameter dimensionality. The configuration combining the low pressure efficiencies increases the computational time, and due to its similar prediction accuracy, it is not seen as at advantage to the original full parameter CEM configuration.

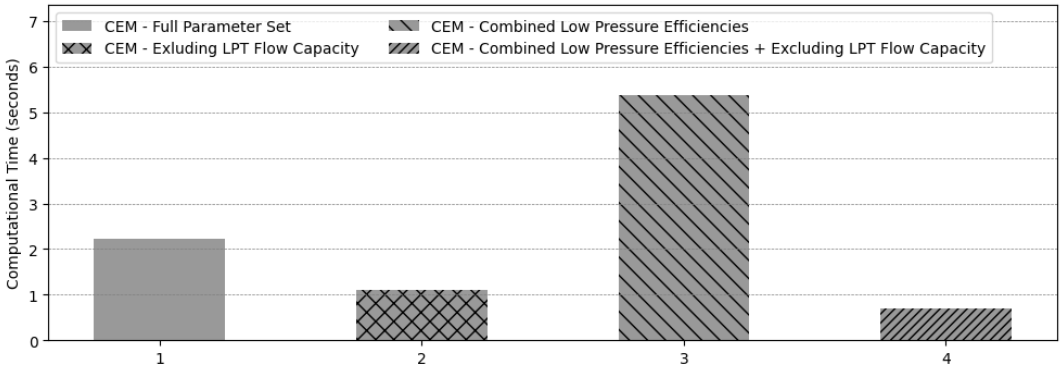

Figure 8.13: Time to run the CEM method per point

8.7. Conclusions on Verification Results

The verification study on potential extensions to the CEM framework demonstrates that the proposed modifications yield prediction performance comparable to that of the original full-parameter configuration.

These alternative configurations rely on several simplifying assumptions. In particular, neglecting deterioration in the LPT flow capacity contradicts the primary objective of this research, which is to estimate all eight health indicators. Although the LPT flow capacity has a weight close to zero for all subsets, these subsets do have a weight regarding the other health indicators. Therefore, these subsets do contribute to the other health indicators and removing them would reduce the diagnostic robustness of the method. In addition, the combined fan core + LPC and LPT efficiency approach assumes identical

deterioration for these components, an assumption that may not be true in practice. Therefore this assumption is not further pursued.

The computational time of the original CEM configuration is fast compared to the current methods used in the industry. Another method to solve the underdetermined problem is using a genetic algorithm. However, this takes approximately 1 hour to solve 1 point [3]. The CEM only takes 2 seconds, which makes it a more than a factor 10^3 faster.

For these reasons, the original full-parameter CEM configuration is kept. This configuration avoids assumptions while maintaining a high level of prediction accuracy. In addition, it demonstrates a high computational efficiency, comparable to or exceeding that of the other proposed configurations.

9

Validation with On-Wing Data

The previous chapter presented the results of the Component Exclusion Method with simulated deterioration. This chapter the CEM will be tested using on-wing data. Three validation studies are presented, all testing the CEM on a different diagnostic feature. First, the data processing of the on-wing data is explained, followed by the three studies. A short background is given on each engine, including the EGT Hot Day Margin, and finally the CEM results of deterioration for each component are presented.

9.1. Data Processing

Prior to applying Adaptive Modeling (AM), the on-wing engine data must undergo extensive preprocessing to ensure meaningful comparison with the GSPy model outputs. A major challenge in this process is the absence of a true *healthy* reference engine, as all engines in operation exhibit some degree of degradation. Furthermore, engine-to-engine variability introduces additional complexity, since individual engines operate with slightly different performance characteristics even under comparable operating conditions. Consequently, careful data processing is required before the data can be used for diagnostic analysis.

9.1.1. Filtering

The on-wing data used in this study is derived from Quick Access Recorder (QAR) measurements. However, this raw data requires substantial preprocessing to ensure consistency, reliability, and suitability for AM simulations. The data extraction and preprocessing procedure consists of the following steps:

1. Account for changes in parameter naming conventions over time. Since parameter names have evolved between the initial data collection period in 2018 and more recent datasets, appropriate filtering is applied to ensure consistent parameter selection across all timestamps.
2. Restrict the dataset to the takeoff, climb, and cruise flight phases, as these phases are most relevant for performance analysis.
3. Apply a strict filter on the rate of change of the core speed, limiting it to a rate of change between -0.01 and 0.01%N2/sec. This ensures that only quasi-steady operating conditions are retained.
4. Filter on N1K between 93% and 101%, the DP is at 97%. Filtering the data to operating points close to the design condition reduces nonlinear deviations between the engine model and actual engine behavior.
5. Remove outliers in secondary performance parameters by applying an Inter Quartile Range (IQR) filter.
6. Exclude climb segments that occur after the first cruise segment to ensure consistency in flight phase progression.
7. Downsample the dataset by selecting operating points closest to the design point N1K. For each flight phase, a maximum of two representative data points per flight per engine is retained, resulting in up to six data points per flight per engine.

After filtering, the dataset presented in Figure 9.1 is kept.

The extracted dataset initially contains combined measurements for both the left and right engines of each aircraft. Therefore, the dataset is first separated into two individual dataframes corresponding to each engine. Subsequently, maintenance records are used to establish engine pairing for each flight. This ensures that the correct engine is tracked consistently over time and allows for accurate analysis of engine-specific performance trends.

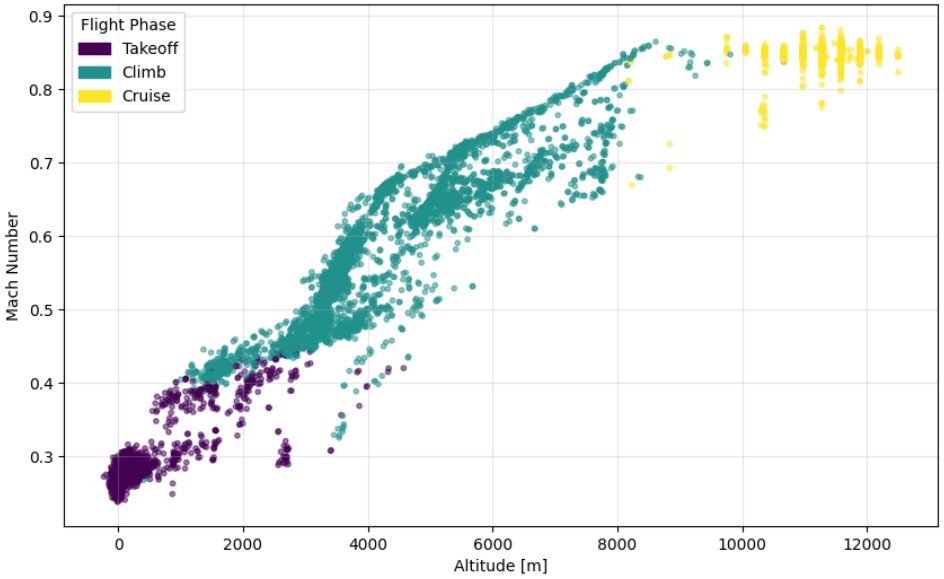

Figure 9.1: Flight phase visualization for QAR flights after filtering

9.1.2. Baselining

Due to discrepancies between GSPy and on-wing data, the Adaptive Modeling tool in GSPy exhibits non-zero predictions from an engine newly installed on-wing. These discrepancies originate from modeling inaccuracies. For instance, the model is created using 1 engine, but other engines may behave slightly differently, causing non-zero deterioration to be present even if the engine is healthy. Furthermore, the on-wing data used in these examples was produced long after the test cell data. To combat this, the results from the first 10 flights after maintenance, waterwash or installation are used to *baseline* the data points that follow. The average HI of these 10 flights is taken and the baselined health indicators are computed using Equation 9.1.

$$HI_{baselined,i} = HI_i - \bar{HI}_{[1:10]} \tag{9.1}$$

This approach effectively removes the GSPy-on-wing offset and ensures that the resulting health indicators represent relative degradation with respect to the assumed healthy engine state, rather than absolute values. In the following sections this is used to show health indicator trends over time.

9.1.3. Companion Engine Analysis

To enhance the interpretation of the health indicator trends, a companion engine analysis is employed. This approach compares the health indicator predictions between the left and right (main and companion) engines under the assumption that both engines operate under similar operating conditions. Given that the engines operate on the same aircraft and are subject to nearly identical flight conditions, systematic differences in the estimated health indicators can be primarily attributed to engine-specific deterioration rather than variations in operational conditions [7]. The main assumptions regarding companion engine analysis have been validated in Appendix D. This analysis therefore provides additional observability that is not available from single-engine evaluation alone.

9.2. HPT Failure Prediction

For each case study approximately 2,000 data points were used for the Component Exclusion Method. The first validation case study concerns an HPT failure in the right engine of a KLM aircraft. This engine was flagged following a Prognos vibrations alert. This led to a request of a borescope inspection at the next arrival. The borescope inspection revealed a fractured HPT stage 1 blade, shown in Figure 9.2. This failure may have occurred as a sudden event or as a result of progressive degradation over time. Therefore, using the CEM, this case study aims to investigate whether indications of prior deterioration, such as erosion, were detectable in the HPT before the failure was detected.

Figure 9.2: HPT blade failure

9.2.1. EGT Hot Day Margin History

The Exhaust Gas Temperature Hot Day Margin (EGT HDM) is a key performance parameter used to assess the health state of the engine. It represents the exhaust gas temperature (T49) margin corrected for hot-day operating conditions and is provided by the engine manufacturer (GE). The evolution of the EGT HDM over time is shown in Figure 9.3. The EGT HDM is consistently lower compared to other engines, and this led to the engine being added to a watch list following an EGT exceedance in August 2023. Such reduced margin is typically due to an already degraded engine, as deterioration increases the required exhaust gas temperature to produce the same thrust. Notably, an increase in EGT HDM is observed shortly before the failure identification. This rise is due to a waterwash performed during that period, which temporary improved engine performance by removing dirt deposits.

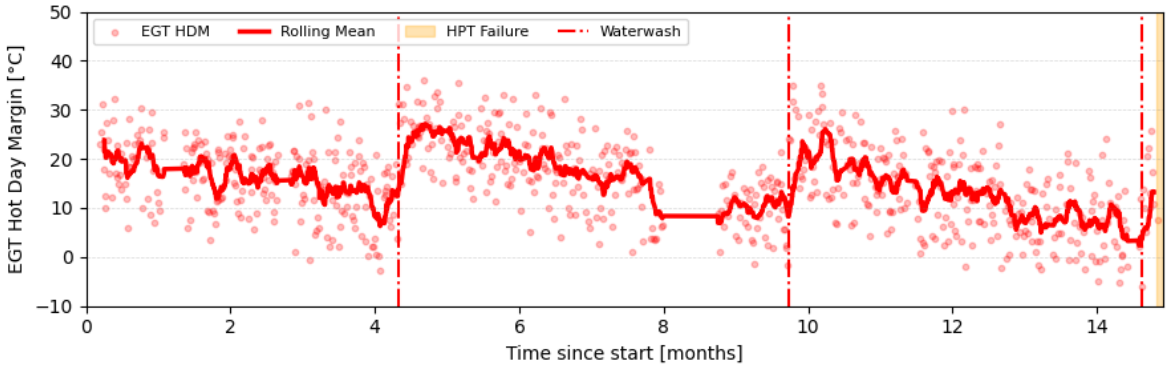

Figure 9.3: EGT HDM before HPT failure

9.2.2. HPT Results

The subset results over time of the engine with the HPT failure are shown in Appendix E. It should be noted that the subset results are before implementing the weight-factor matrix. The results of the CEM

analysis are presented in Figure 9.4 and Figure 9.5. The two companion engines are baselined with respect to the start of their pairing period. Over the analysed period, two waterwashes were performed on engine 2. The effect of waterwashes on HPT performance will be presented in subsection 9.4.3

The rolling average of the HPT efficiency reveals a clear divergence between the two engines. While the HPT efficiency health indicator of engine 1 remains approximately constant around 0%, engine 2 exhibits a gradual deterioration reaching values of approximately -1%. Although such a level of degradation may fall within normal degradation patterns, the absence of a similar trend in the companion engine suggests an abnormal behavior. This discrepancy may therefore be indicative of early-stage deterioration in the HPT of engine 2.

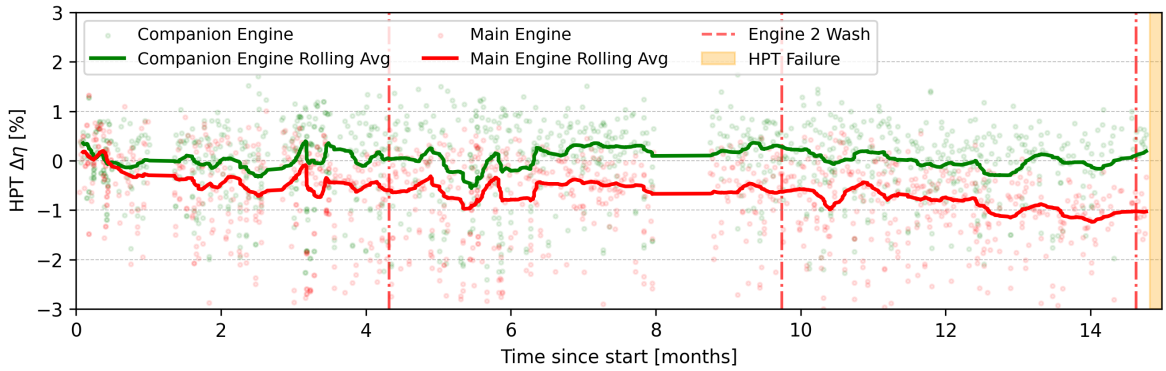

Figure 9.4: HPT efficiency

The HPT flow capacity health indicators of engines 1 and 2 exhibit a similar trend up to approximately August 2023, with both engines showing a gradual increase consistent with normal degradation behavior. However, from this point onward, the flow capacity of engine 2 begins to increase at a higher rate compared to engine 1. While such behavior may not be immediately apparent when analyzing a single engine in isolation, it becomes evident through direct comparison with its companion engine.

This observation highlights the value of companion engine analysis, where relative deviations can reveal subtle anomalies that would otherwise remain undetected. In particular, the use of residuals provides a more sensitive means of identifying divergence in engine condition, making it a useful diagnostic tool.

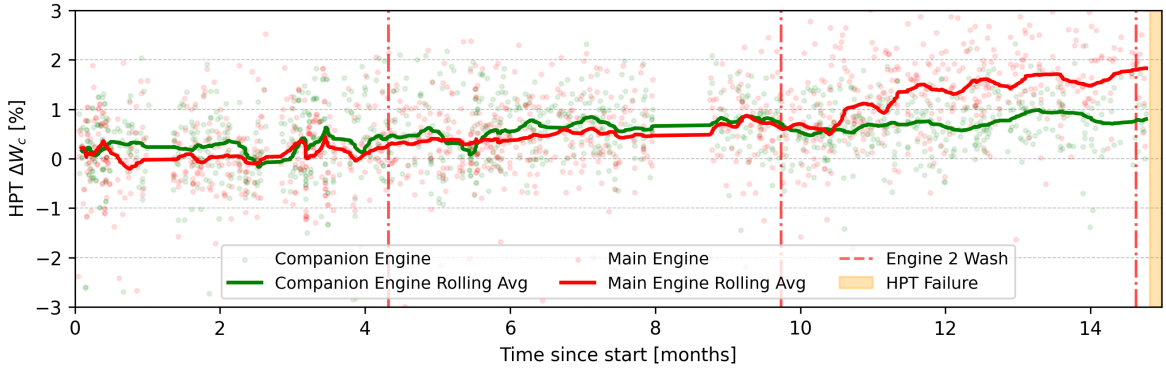

Figure 9.5: HPT flow capacity

9.2.3. Residuals

Residual are stated as the difference between the health indicator value of the left (engine 1) and right (engine 2) engines. In this study, the right engine (engine 2) is assumed to be more deteriorated than the left engine (engine 1). Therefore, the residual is defined as:

$$R_{hi}(t) = X^2_{hi}(t) - X^1_{hi}(t) \tag{9.2}$$

The residuals of the HPT efficiency and flow capacity are presented in Figure 9.6. In this figure, regions are classified as normal or abnormal based on whether the residual of both health indicators of one component, in this case the HPT efficiency and HPT flow capacity, exceed a predefined threshold of 0.5% or -0.5% for a sustained period of 10 data points.

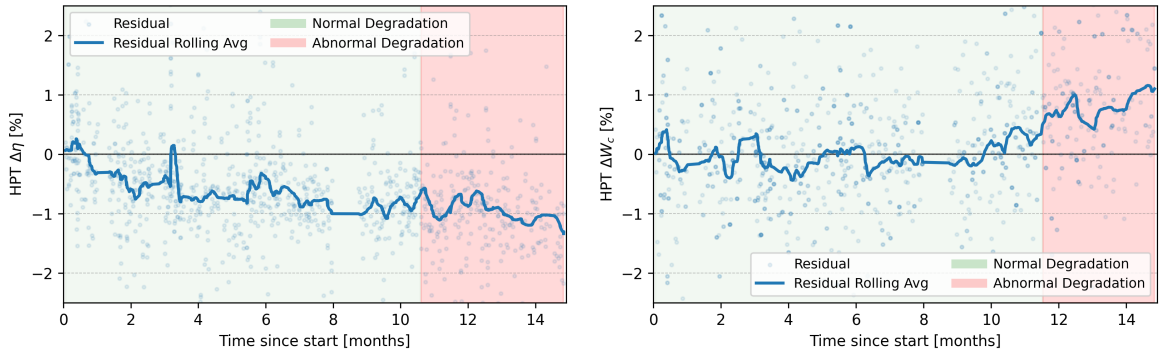

(a) HPT efficiency

(b) HPT flow capacity

Figure 9.6: Residuals of HPT health indicators (right HIs - left HIs)

The HPT efficiency residual becomes abnormal relatively early in the observation period. While this may indicate early-stage HPT deterioration in engine 2, alternative explanations should also be considered, such as differences in engine age or prior operational history. Older engines may exhibit distinct degradation patterns compared to newer or recently overhauled counterparts. This limitation is partly accounted for using the baselining method.

The HPT flow capacity residual also indicates abnormal behavior, but with a more sudden shift. For the majority of the observation period, the residual fluctuates around zero, suggesting comparable degradation between both engines. However, approximately three months prior to failure detection, a distinct increase in the residual is observed. Such a rise in flow capacity is consistent with structural damage, such as the fracture or partial loss of a turbine blade, which effectively increases the flow area through the turbine.

9.2.4. Conclusion

The results demonstrate that indications of the HPT failure were present well before its identification during the borescope inspection. This is most clearly observed in the HPT flow capacity residual, which exhibits a distinct deviation approximately three months prior to the inspection. When considering only the rolling average of the health indicators for engine 2, this behavior could be interpreted as normal degradation. However, the residual-based comparison with the companion engine reveals that the observed trend is in fact anomalous.

This case highlights the added value of residual-based companion engine analysis in enhancing fault detectability. By leveraging relative deviations rather than absolute trends, the method enables earlier and more reliable identification of abnormal behavior. Consequently, residual analysis proves to be a powerful tool for condition monitoring, offering improved sensitivity to faults and supporting more timely maintenance decision-making.

9.3. Performance Restoration Shop Visit

The next validation case study concerns a HPT repair on engine 956XXX of the GEnx-1B. The objective of this case study is to assess whether the recovery of engine health indicators following an overhaul can be captured using the CEM framework.

The engine was removed from service following a Prognos alert, after which a full Performance Restoration Shop Visit (PRSV) was carried out. This particular PRSV involved the overhaul of both the combustor and the HPT, which is typically scheduled after approximately 3,000 operating cycles. A preceding borescope inspection identified tip damage on the stage 2 HPT blades. Further inspection revealed the presence of shroud droop. The shroud, forming a circumferential ring around the turbine blades,

serves to minimize tip leakage and guide cooling flow. In the event of shroud droop, the casing deforms radially inward, reducing clearance and potentially leading to contact between the blade tips and the shroud. This interaction accelerates wear on both components, resulting in a reduction in turbine efficiency and an increase in flow capacity. Based on these findings, the replacement of the HPT stage 2 blades was combined with the scheduled PRSV.

9.3.1. EGT Hot Day Margin History

The evolution of the EGT HDM over time is presented in Figure 9.7. The vertical dashed lines indicate the occurrence of water washes for the engine considered in this case. The impact of these maintenance actions is clearly reflected in the EGT HDM, which shows a noticeable increase following each water wash due to the temporary restoration of component performance. In contrast, a decline in EGT HDM is observed leading up to the shop visit, ultimately triggering a Prognos alert. This reduction in EGT HDM is consistent with progressive engine degradation and serves as a key indicator for maintenance planning.

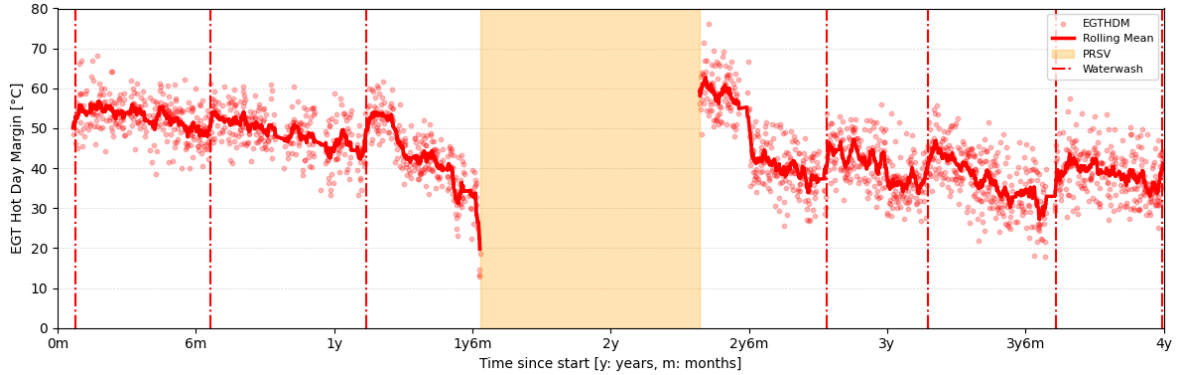

Figure 9.7: EGT HDM history

9.3.2. HPT Results

Figure 9.8 shows the HPT efficiency deterioration predicted by the Component Exclusion Method. Given that this case involves HPT damage, a reduction in efficiency prior to the shop visit is expected and is indeed observed. The HPT efficiency decreases from approximately 0% to -1% over a period of around four months. A similar degradation trend is visible in the companion engine, which underwent a similar workscope approximately one month earlier. Following the performance restoration shop visit, a significant recovery in HPT efficiency is observed, with an increase of nearly 2%, indicating successful restoration of component performance.

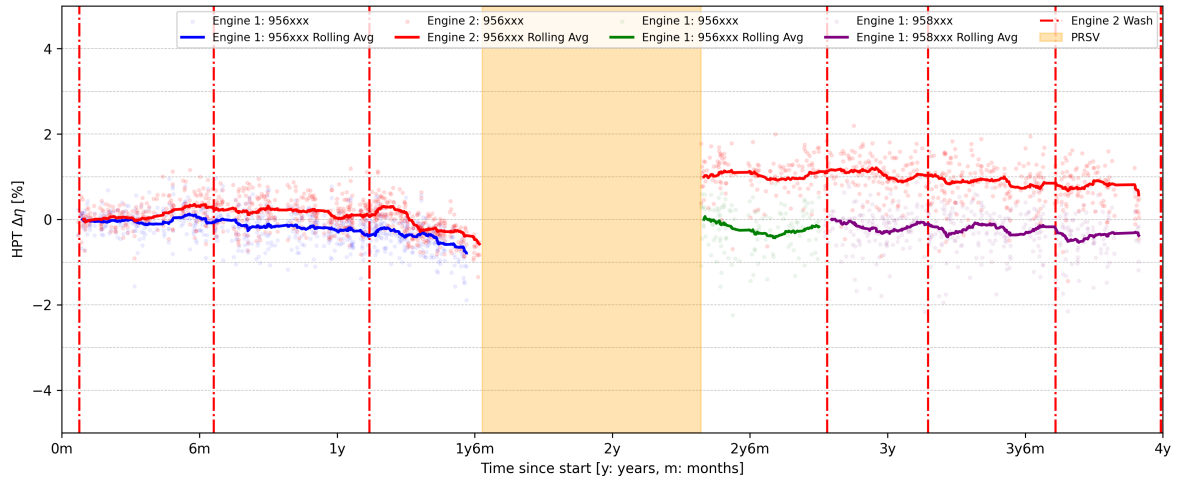

Figure 9.8: HPT efficiency

The HPT flow capacity results are shown in Figure 9.9. In contrast to efficiency, deterioration of the HPT is associated with an increase in flow capacity. This trend is clearly evident, as the flow capacity gradually rises from approximately 0% to 1% prior to the shop visit. After the restoration, the flow capacity returns to values close to the baseline, which confirms the maintenance action had an effect on the performance of the HPT.

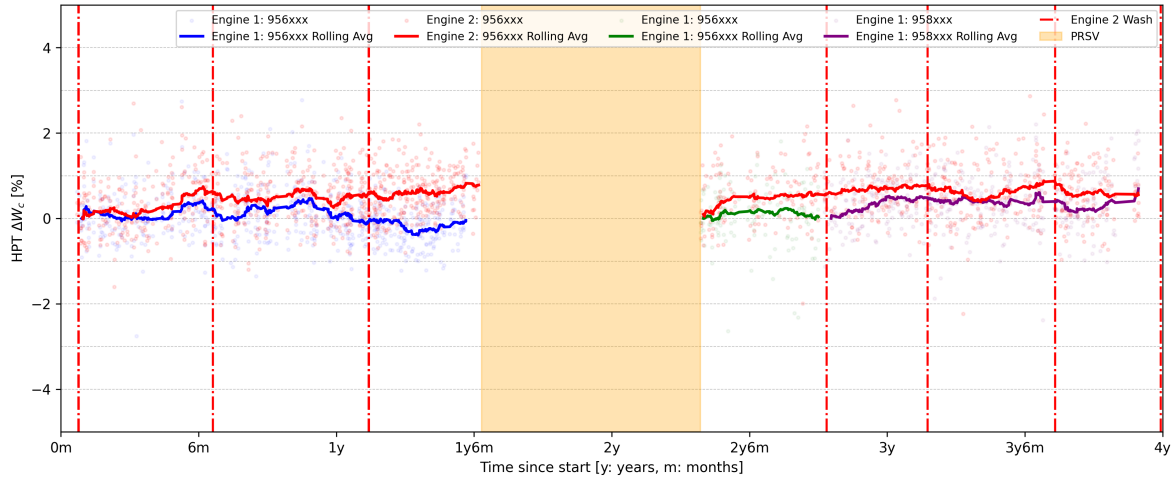

Figure 9.9: HPT flow capacity

9.3.3. Conclusion

The results of this case study demonstrate that the Component Exclusion Method (CEM) is capable of effectively capturing both HPT deterioration and subsequent performance recovery. Importantly, this also confirms findings from the previous case study, where the method already demonstrated the ability to reliably predict impending HPT failure well in advance of the actual event. Therefore, this case confirms two important characteristics of the CEM.

Firstly, the predicted deterioration trends in both HPT efficiency and flow capacity are consistent with inspection findings, indicating that the method accurately captures the underlying physical degradation processes leading up to failure. This reinforces the earlier conclusion that CEM is not only sensitive to gradual performance loss, but also capable of identifying deterioration patterns that precede HPT-related incidents.

Secondly, following the Performance Restoration Shop Visit, both health indicators show a clear recovery, indicating increased efficiency and a reduction in flow capacity deviation consistent with the expected effects of HPT maintenance. This post-maintenance behavior further validates the physical interpretability of the method.

9.4. Waterwashes

The second case study investigates variations in engine performance due to waterwashes. A waterwash is a maintenance procedure used to remove fouling deposits present in compressors. As illustrated in Figure 9.7, a waterwash typically increases the EGT HDM by around 10 degrees, indicating a temporary recovery of engine performance.

Due to the absence of suitable failure cases in the available dataset, the capability of the Component Exclusion Method (CEM) to detect deterioration in the fan core + LPC and HPC cannot be directly validated using fault events. Instead, waterwashes are used to assess the sensitivity of the method to changes in the condition of these components.

To better capture changes in performance trends before and after waterwash event, the rolling mean of each health indicator is reset following a waterwash. In total, two cases are selected. These cases corresponds to aircraft with engines with long operational periods and multiple waterwash event. For clarity, these are referred to as aircraft A and aircraft B. The HPT predictions are also included, although these are not expected to change due to waterwashes. The LPT is not included, as the effects of

waterwashes to the LPT health indicators are similar or less than the HPT health indicators.

9.4.1. Fan core + LPC Results

The results of the fan core + LPC efficiency are presented in Figure 9.10. The impact of waterwashes on this health indicator appears limited. Small variations in the rolling average are visible, the overall efficiency remains close to zero, which is consistent with earlier verification results, which showed that the fan core + LPC efficiency is predicted with relatively low accuracy by the subsets, and therefore it is given a low weight in the weight-factor matrix. This causes the value to stay close to zero.

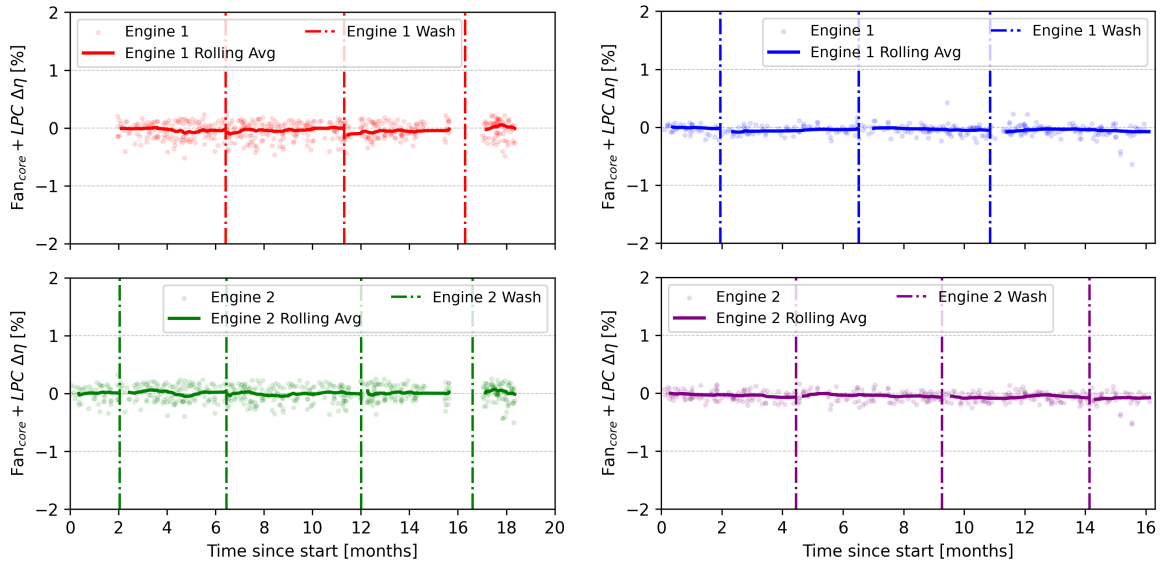

(a) Example of aircraft A

(b) Example of aircraft B

Figure 9.10: Effect of waterwashes on the fan core + LPC efficiency

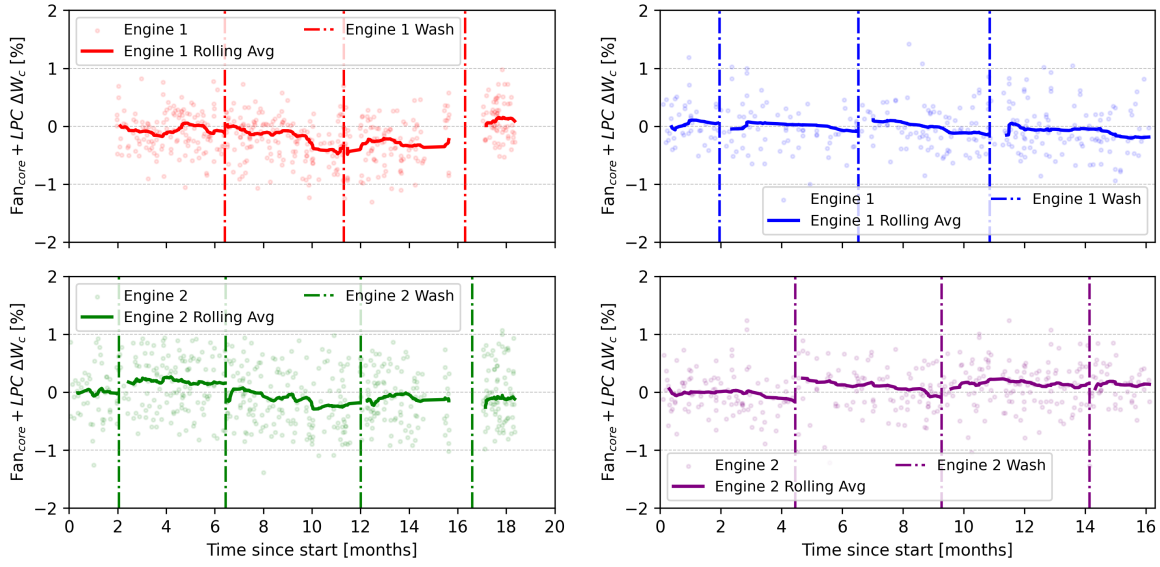

(a) Example of aircraft A

(b) Example of aircraft B

Figure 9.11: Effect of waterwashes on the fan core + LPC flow capacity

The fan core + LPC flow capacity results are shown in Figure 9.11. In contrast to the efficiency, a more noticeable response to waterwash events is observed. Specifically, a slight upward shift in flow

capacity occurs following each waterwash, consistent with the expected removal of dirt. In addition, a gradual downward trend is visible over time, reflecting the slow deterioration of the fan core + LPC.

9.4.2. HPC Results

The results for HPC efficiency and flow capacity are presented in Figure 9.12 and Figure 9.13, respectively. In contrast to the fan core and LPC, the HPC health indicators exhibit a pronounced response to waterwash events, with all cases showing a clear and immediate recovery in performance following each intervention.

HPC efficiency increases by approximately 1–2% after each waterwash, while flow capacity improves by approximately 0.5–1% on average. In addition, both health indicators exhibit a persistent downward trend over time. This gradual degradation, which is repeatedly reversed by waterwash events, is consistent with fouling-related deterioration.

As summarized in Table 2.1, fouling refers to the accumulation of contaminants on compressor or turbine surfaces, leading to increased surface roughness and reduced aerodynamic performance. In the HPC, this manifests as a reduction in both efficiency and mass flow capacity due to increased losses and altered flow conditions. The overall deterioration trend, combined with the performance recovery following waterwash events, supports the conclusion that fouling is a dominant degradation mechanism in this component, and that this is correctly predicted by the CEM.

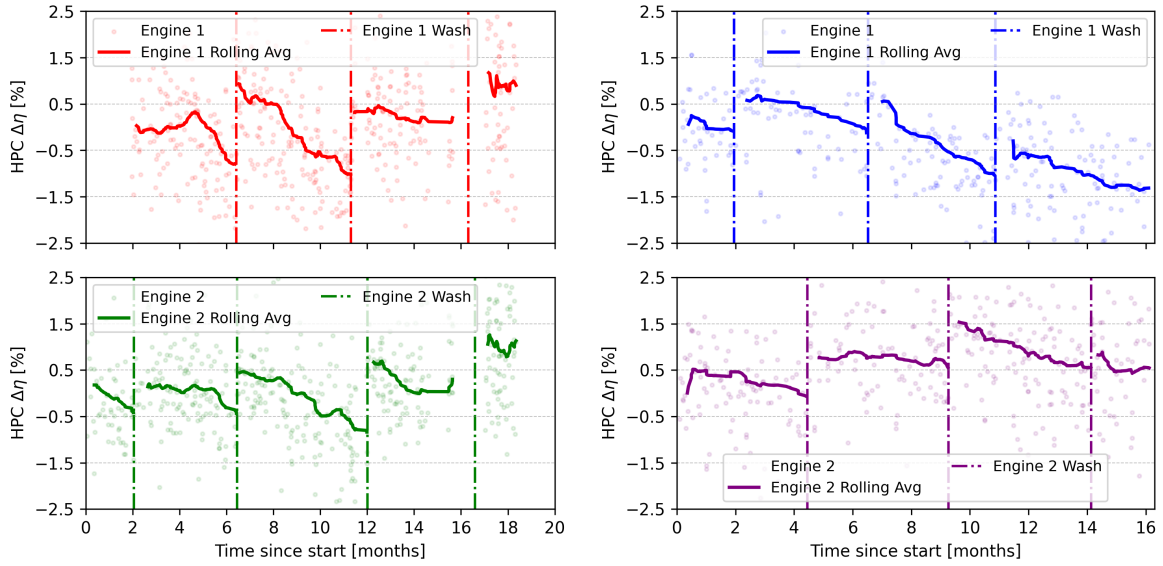

(a) Example of aircraft A

(b) Example of aircraft B

Figure 9.12: Effect of waterwashes on the HPC efficiency

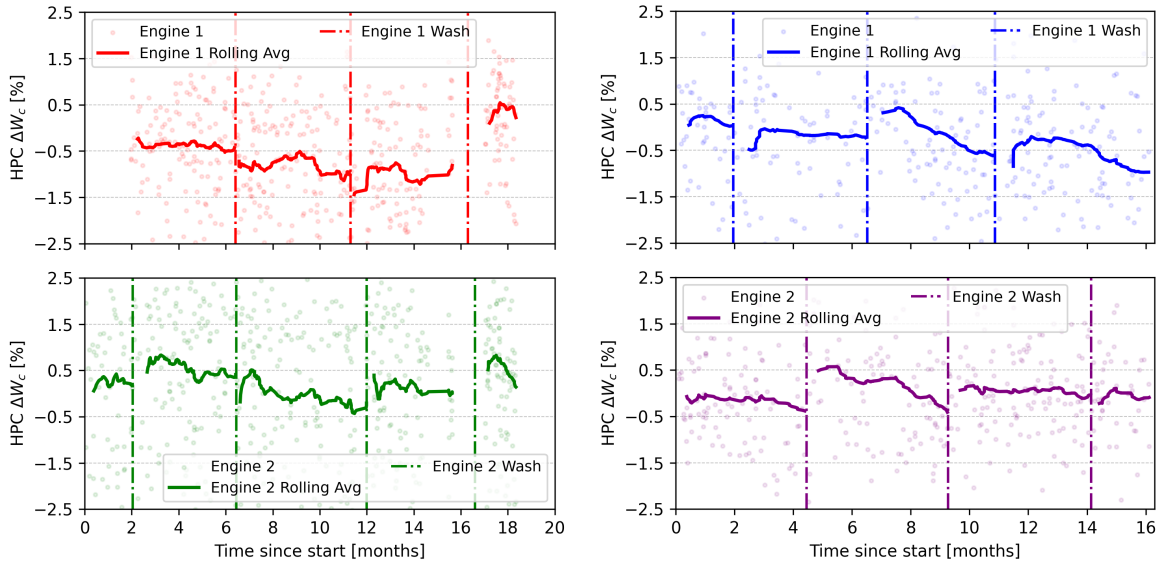

(a) Example of aircraft A

(b) Example of aircraft B

Figure 9.13: Effect of waterwashes on the HPC flow capacity

9.4.3. HPT Results

The effect of waterwashes for the HPT health indicator predictions are presented below. As expected, the effect of waterwashes is limited. In most cases, the efficiency remains constant, with only minor increases observed in some cases. This is consistent with the fact that waterwashes mainly target compressor fouling.

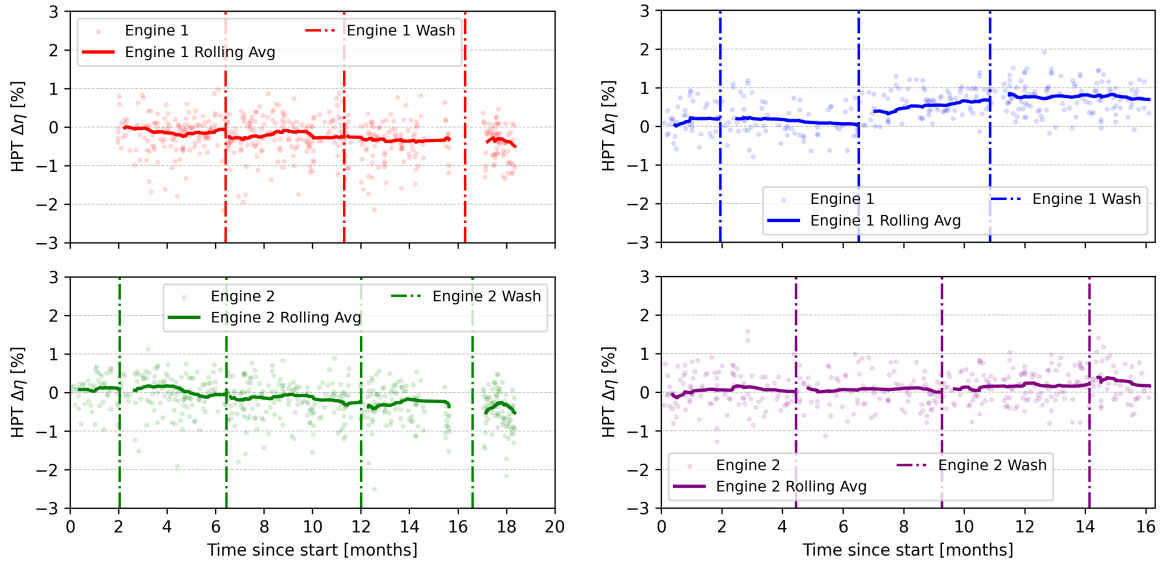

(a) Example of aircraft A

(b) Example of aircraft B

Figure 9.14: Effect of waterwashes on the HPT efficiency

In addition, the HPT flow capacity also shows a limited sensitivity to waterwashes. In some cases a decrease of flow capacity is seen after a waterwash, before quickly returning to the original state. This is contradictory with expectations, as the removal of residue should increase the HPT flow capacity. This phenomenon could be due to smearing from the HPC to the HPT.

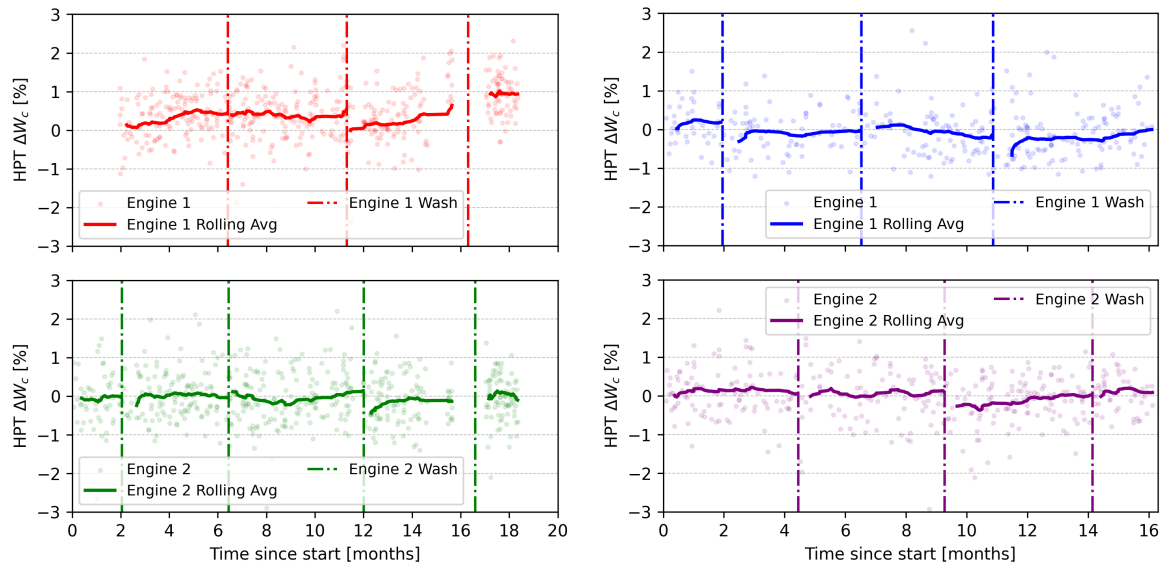

(a) Example of aircraft A

(b) Example of aircraft B

Figure 9.15: Effect of waterwashes on the HPT flow capacity

9.4.4. Conclusion

The waterwashes are visible in both fan core + LPC and HPC predictions. The fan core + LPC health indicators show only limited responses to waterwashes, particularly the efficiency, which remains close to zero. In contrast, the HPC health indicators display clear improvements after each waterwash. The findings therefore suggest that the CEM is capable of capturing performance changes in the fan core + LPC and HPC due to fouling. Furthermore, the waterwash has negligible effects on HPT performance, which is in line with expectations.

9.5. Discussion

The goal of this validation study was to assess whether the Component Exclusion Method (CEM) is capable of predicting component deterioration. Across the three cases analyzed in this chapter, the results indicate that this is indeed the case, although with varying levels of certainty and robustness depending on the scenario.

In Case 1 (HPT failure prediction), the CEM demonstrates strong predictive capability. Firstly, the method is able to indicate the onset of HPT deterioration prior to detection by other monitoring systems, suggesting an advantage in early fault identification. Secondly, when combined with companion engine analysis, the consistency of the observed degradation trends is improved, reinforcing the interpretability of the predicted health indicators. However, a limitation is observed in the form of a sudden drop in the rolling average of HPT efficiency of approximately −1%, which may indicate a CEM discrepancy or sensitivity to variations in operating points. Since the true degradation is not directly observable over time, it remains difficult to fully validate whether this represents actual physical behavior or a modeling discrepancy.

In Case 2 (Performance Restoration Shop Visit), the CEM again shows two clear strengths. Firstly, the method captures gradual deterioration trends, indicating prediction robustness of the CEM. Secondly, the spread of estimated health indicator values remains relatively low, suggesting that the CEM is able to produce stable predictions of component behavior under near-steady-state conditions. On the downside, the severity of the EGT HDM drop is not represented by CEM. While the EGT HDM drop very sudden near the failure region, this is not visible in the CEM plot, where the last points represent a -1% HPT efficiency deterioration and a 2% flow capacity deterioration. This presents a limitation in the fact that it is uncertain if CEM can accurately quantify sudden deterioration.

In Case 3 (Waterwash response analysis), the CEM demonstrates a mixed but physically largely consistent response to short-term maintenance actions. First, the method is able to clearly capture per-

formance recovery in the HPC health indicators following each waterwash event, with both efficiency and flow capacity exhibiting distinct and repeatable improvements. This indicates that the CEM is not only sensitive to long-term degradation trends, but also to relatively small, maintenance-induced shifts in compressor condition. Moreover, the consistency of this response across multiple waterwash events and across all four analyzed engines suggests that the method is robust in detecting repeatable performance changes under varying operating histories.

Importantly, the observed behavior is physically consistent with fouling-related deterioration mechanisms. The clear downward trend in HPC efficiency and flow capacity prior to each waterwash, followed by a stepwise recovery after cleaning, confirms that the model is effectively capturing fouling accumulation and removal. As fouling primarily affects compressor aerodynamics through surface contamination and increased losses, the CEM's ability to track these reversible changes demonstrates that it can reliably identify HPC deterioration driven by fouling mechanisms.

However, a limitation is observed in the fan core + LPC and HPT results, where the expected health indicator response is either weak or inconsistent. In particular, the near-zero sensitivity of the fan core + LPC efficiency and the inconsistent behavior in HPT flow capacity (including small decreases following waterwashes) indicate potential cross-component smearing effects or reduced observability of these components under the current model formulation.

Finally, it is worth noting that the spread of health indicators remains low throughout the validation cases. This indicates a high level of stability and robustness in the health indicator predictions. The limited spread may be a consequence of the novel filtering method, which enhanced the homogeneity of the dataset. Alternatively, it may also be a characteristic of the Component Exclusion Method, which produces smoother health indicators predictions due to its modeling structure.

10

Conclusion

Condition-based maintenance is increasingly considered a promising approach for maintenance scheduling [2, 29, 27]. While a variety of methods have been proposed for engine health monitoring and prognostics, their operational implementation remains limited due to uncertainties in modeling approaches, sensor limitations, and robustness under real-world conditions.

The research question at the heart of this study was: *How can the Component Exclusion Method and Companion Engine Analysis be used to improve condition monitoring accuracy of the GEnx-1B with fewer sensors?*

Following the selection of the gas path analysis tool GSPy, a methodology was developed to construct the Component Exclusion Method (CEM), which combines health indicator prediction of multiple subsets into a single prediction using a weight-factor matrix. The method was subsequently verified and validated using both simulated deterioration cases and on-wing data. A number of conclusions can be drawn from the research, both in regards to the modeling method and the final results. These conclusions are presented below.

- **Existing modeling approaches:** To address the limited sensor availability in the GEnx-1B engine, several gas path modeling approaches have been proposed in the literature. These methods typically aim to predict component health parameters using data-driven or hybrid methods such as genetic algorithms [3, 7, 6, 24, 5]. Despite their theoretical capabilities in handling underdetermined systems and nonlinear engine behavior, these approaches have seen limited adoption in industry, primarily due to their relatively high computational cost when applied to large-scale, high-frequency datasets. A key advantage of the CEM lies in its computational efficiency. Whereas alternative methods may require approximately an hour to estimate health parameters for a single operating point, the CEM operates deterministically and evaluates a single subset in approximately 0.3 seconds. Since six subsets are being used and accounting for some post-processing time, this results in a total processing time of roughly 2-5 seconds per data point. This level of performance enables near real-time applicability, where incoming flight data can first be filtered for steady-state conditions, after which the CEM can be deployed for on-wing health indicator estimation.
- **QAR data:** The use of QAR data enabled the application of advanced filtering techniques, significantly improving the quality and consistency of the dataset. By utilizing 1 Hz measurements, operating points could be filtered on near steady-state conditions, secondary performance parameters, and operating conditions close to the design point. This resulted in a more homogeneous dataset and limited the spread of the health indicator predictions. A limitation of the QAR dataset was that the available data only started from mid-2020 onward, whereas the limited number of available HPC failure cases occurred prior to this period. Nevertheless, QAR data proved to be significantly more suitable than snapshot data, as snapshot data does not allow filtering on steady-state conditions, secondary performance parameters, or proximity to the design point through corrected N1 values. Furthermore, as shown in Appendix D, the companion engine validation demonstrated that companion engines operate under highly similar input conditions, thereby im-

proving the reliability and accuracy of companion engine analysis.

- **GSPy:** GSPy has proven to be a strong tool for implementing health indicator prediction using on-wing engine data, particularly due to its flexibility in integrating additional customizations and extensions required for this specific research. Measurement tolerances were included to account for sensor inaccuracies, which is a novel feature not implemented in GPA before. In addition, the crossflow was implemented in GSPy during this thesis, which improved model matching to on-wing data. This model demonstrated excellent computational efficiency, making it suitable for simulation on large on-wing datasets.
- **Health Indicator Subsets:** Due to a limited number of sensors available, adaptive modeling was performed using various subsets of health indicators. This led to varying results, from non-converging subsets, to very accurate subsets, to poorly predicting subsets. Although these subsets could not be directly applied, their behavior revealed a promising concept for health indicator estimation. By combining the best-performing health indicators from each subset, it was possible to make an accurate prediction.
- **CEM Verification with simulated deterioration:** The CEM verification results with simulated data demonstrated strong predictive performance. Especially the core components, the HPC and HPT, could be predicted with great accuracy. The low-pressure components showed higher degrees of uncertainty, especially the fan core + LPC efficiency and the LPT flow capacity. Furthermore, the extensions to the CEM showed advantages and limitations. Although the prediction performance was similar to the original CEM configuration, disregarding health indicators contradicts the primary objective of this research. In addition, the most notable improvement was the computational efficiency, and this was already deemed fine with the original CEM configuration of 6 subsets.
- **CEM Validation with on-wing data:** The CEM has been validated with on-wing data. The first validation case consisted of a stage 1 HPT blade failure, where after a borescope inspection a broken off blade was found. Using historical on-wing data the CEM was able to construct trends over time for all 8 health indicators, and could identify abnormal deterioration trends before existing monitoring alerts were triggered.

  In addition, a strong recovery of HPT health indicators was seen following a Performance Restoration Shop Visit (PRSV), which consisted of maintenance of the combustor and HPT. Both the HPT efficiency and flow capacity improved significantly after the shop visit, which indicates the robustness of the CEM not only to deterioration, but also performance improvements.

  Furthermore, due to the absence of HPC failure cases, waterwashes were used to validate the CEM on compressor performance. The HPC was predicted accurately during waterwashes, where the efficiency and flow capacity health indicators improved during a waterwash. However, the fan core + LPC efficiency showed limited improvement.
- **Companion Engine Analysis:** The Companion Engine Analysis (CEA) improved diagnostic fidelity by enabling the identification of abnormal deterioration between companion engines. By analyzing the residuals of health indicators between companion engines, discrepancies in health indicators were made visible, allowing a distinction to be made between normal degradation trends and abnormal behavior.

  Using engineering judgment, threshold values were defined for the residuals of the health indicators. When both indicators exceeded these thresholds for a sustained period, this degradation of this component is classified as abnormal degradation. The first validation case demonstrated the added diagnostic value of this method, as it enabled earlier and clearer identification of abnormal engine behavior, later resulting in a failure.

This work is novel because it restructures the underdetermined GPA problem into a set of solvable subproblems and combines the solutions of those subproblems using weight-factors, thereby avoiding the computational cost of optimization-based approaches. Furthermore, the integration of Companion Engine Analysis introduces a residual-based diagnostic layer that enables the detection of abnormal degradation. Together, these two contributions were demonstrated on the GEnx-1B under operational on-wing conditions.

11

Recommendations

Although the CEM with Companion Engine Analysis showed promising results, work can still be done to improve the method. The following recommendations can improve the framework, both increasing prediction accuracy and computational efficiency.

- A key recommendation for future GPA research is to improve the fidelity of the engine model. Currently, the model is based on test-cell data from a single engine, which introduces bias when adaptive modeling is applied to other engines and negatively affects prediction accuracy. Developing engine-specific performance maps would likely improve model fidelity and diagnostic performance. One potential approach is to use neural networks trained on *healthy* on-wing data from individual engines to generate engine-specific performance maps.
- In this study, the health indicators of only four components were used. During GSPy verification, a model was created in which the fan core and LPC were separated into two components, thereby creating a 5th component. Preliminary investigations showed the increased model complexity introduced by separating the fan core and LPC did not show significant improvements to justify further investigation within the scope of this research. However, this could improve modeling accuracy. In particular, the current implementation of the crossflow with a Cf factor of 0.5 is now physically inconsistent, as airflow cannot flow back to the fan duct after the LPC. Revisiting this model configuration may help increase prediction accuracy, and remove physical inconsistency.
- Preliminary testing of the subsets found that only 12 out of the total 28 subsets converge under full deterioration conditions. Appendix C provides an overview of the non-converging subsets. While part of these subsets can be explained analytically, such as the necessity of including the HPT flow capacity, the remaining non-converging subsets have not been fully derived mathematically. Instead, a consistent trend was identified: both the compressor side (fan core + LPC and HPC) and the turbine side (HPT and LPT) require at least three out of their four corresponding health indicators to ensure convergence. Establishing a mathematical explanation for all non-converging subsets would be valuable for future work, as it would improve the understanding of the underlying behavior of Gas Path Analysis (GPA) and provide a stronger foundation for subset selection.
- To account for measurement uncertainty, tolerances were introduced in the adaptive modeling framework within GSPy. These tolerances provide the solver with increased flexibility when matching measured parameters during the solution of the system of equations. However, their implementation did not result in improved performance when handling noisy measurements. This indicates that further work is required to develop more robust methods for dealing with measurement noise. Additionally, the current implementation does not include tolerances for ambient conditions. Extending the framework to incorporate uncertainties in ambient parameters could further enhance the overall model accuracy.
- Currently, the Component Exclusion Method is trained and tested on the model of a single engine, and using three operating conditions reflecting the takeoff, climb and cruise flight phases. This should be adapted by training the CEM on more operating points. One potential approach is to train a neural network that predicts the optimal weight-factor matrix as a function of operating conditions.

- Besides thrust, the GEnx-1B provides electrical power to electric systems. A certain amount of energy is extracted, called power off-take. GSPy currently does not enable users to specify the power off-take. Since it is measured during flight, it can be incorporated into the adaptive modeling simulation, thereby increasing prediction accuracy.
- During the final stages of this thesis research, a new update to GSPy introduced humidity modeling capabilities [8]. Although humidity effects are expected to have a limited influence under cruise operating conditions, it can significantly affect engine performance during takeoff and other low-altitude operating conditions. Incorporating humidity effects into the adaptive modeling framework could therefore improve prediction accuracy, particularly during these flight phases. Research conducted by van Vuuren demonstrated that ambient humidity conditions around airports can be estimated using local weather station data [50]. Integrating such data into the adaptive modeling framework would provide a practical approach for accounting for humidity variations in on-wing engine performance analysis.
- Finally, the validation of the CEM for the compressors has so far been validated using water-washes. Although this can predict improvements in HPC performance, the CEM also has to be validated for a fan core/LPC or HPC failure case. A waterwash only represents performance recovery rather than performance deterioration. Including these failure cases would improve the assessment of the method to compressor degradation.

References

[1] International Air Transport Association (IATA). *IATA 2025 Safety Report*. Accessed: 2026-04-21. Mar. 2026. URL: https://www.iata.org/en/pressroom/2026-releases/2026-03-09-01/.

[2] A. J. Volponi and L. Tang. "Principles of Gas Path Analysis". In: (2026). DOI: 10.1201/9781003649274.

[3] T. O. Rootliep, W. P. J. Visser, and M. Nollet. "Evolutionary Algorithm for Enhanced Gas Path Analysis in Turbofan Engines". In: *Volume 1: Aircraft Engine; Fans and Blowers; Marine; Wind Energy; Scholar Lecture*. Virtual, Online: American Society of Mechanical Engineers, 2021, V001T01A011. DOI: 10.1115/GT2021-59089.

[4] M. Otten. "Development of a Diagnostics Model for the GEnx-1B Turbofan Engine Using On-Wing Performance Data". MA thesis. Delft University of Technology, 2021.

[5] B. de Bruin. "Surrogate Model Based Diagnostics for the GEnx-1B Turbofan Engine". MA thesis. Delft University of Technology, 2022.

[6] W. Brachmi. "Extended Surrogate Modelling for Gas Turbine Diagnostics & Prognostics". MA thesis. Delft University of Technology, 2023.

[7] A. Singhvi. "Residual-Based AI Model for Turbofan Engine Predictive Maintenance". MA thesis. Delft University of Technology, 2025.

[8] W. Visser. *GSPy: Python Propulsion and Power System Performance Modelling and Simulation Tool*. GitHub repository, version 1.3.0.1 (12-Nov-2025). URL: https://github.com/wvisser1958/GSPy.

[9] H. I. H. Saravanamuttoo, H. Cohen, and G. F. C. Rogers. *Gas Turbine Theory*. 5th. Harlow: Prentice Hall, 2007. ISBN: 978-0-13-015847-5.

[10] NASA. *Turbine Engine Thermodynamic Cycle - Brayton Cycle*. Accessed: 2025-09-23. 2021. URL: https://www.grc.nasa.gov/www/k-12/airplane/brayton.html.

[11] M. Pini. *AE4238: Aero Engine Technology — Turbomachinery: Performance Maps*. Lecture slides, Course material for AE4238. 2022.

[12] L. A. Urban. "Gas Path Analysis Applied to Turbine Engine Condition Monitoring". In: *Journal of Aircraft* 10.7 (1973), pp. 400–406. DOI: 10.2514/3.60240.

[13] G. Bechini. "Performance Diagnostics and Measurement Selection for On-Line Monitoring of Gas Turbine Engines". Supervisor: Professor Riti Singh. PhD thesis. Cranfield University, School of Engineering, Dec. 2007.

[14] P. C. Escher. "Pythia: An Object-Oriented Gas Path Analysis Computer Program for General Applications". Supervisor: R. Singh. PhD thesis. Cranfield University, School of Mechanical Engineering, Oct. 1995.

[15] GE Aerospace. *Air France-KLM Selects GEnx Engines for Boeing 787 Fleet*. Accessed: 2025-09-18. Mar. 2014. URL: https://www.geaerospace.com/news/press-releases/commercial-engines/air-france-klm-selects-genx-engines-boeing-787-fleet.

[16] GE Aviation. *GEnx-1B Turbofan Engine Installation Manual*. Revision 28. GE Proprietary Information, Confidential. May 2024.

[17] F. Forest. "Unsupervised Learning of Data Representations and Cluster Structures: Applications to Large-scale Health Monitoring of Turbofan Aircraft Engines". PhD thesis. Université Paris-Saclay, 2022.

[18] M. Valkier. "Turbofan HPT Blade Failure Prediction Based on Vibration Decomposition: A Real-World Case Study". MA thesis. Delft University of Technology, 2024.

[19] M. G. De Giorgi, N. Menga, and A. Ficarella. “Exploring Prognostic and Diagnostic Techniques for Jet Engine Health Monitoring: A Review of Degradation Mechanisms and Advanced Prediction Strategies”. In: *Energies* 16.6 (2023), p. 2711. DOI: 10.3390/en16062711.

[20] A. J. Volponi. “Gas Turbine Engine Health Management: Past, Present, and Future Trends”. In: *Journal of Engineering for Gas Turbines and Power* 136.5 (2014), p. 051201. DOI: 10.1115/1.4026126.

[21] A. K. S. Jardine, D. Lin, and D. Banjevic. “A Review on Machinery Diagnostics and Prognostics Implementing Condition-Based Maintenance”. In: *Mechanical Systems and Signal Processing* 20.7 (2006), pp. 1483–1510. ISSN: 0888-3270. DOI: 10.1016/j.ymssp.2005.09.012.

[22] K. Khan et al. “Recent Trends and Challenges in Predictive Maintenance of Aircraft’s Engine and Hydraulic System”. In: *Journal of the Brazilian Society of Mechanical Sciences and Engineering* 43.8 (2021). DOI: 10.1007/s40430-021-03121-2.

[23] P. Bonissone, X. Hu, and R. Subbu. “A Systematic PHM Approach for Anomaly Resolution: A Hybrid Neural Fuzzy System for Model Construction”. In: *Annual Conference of the PHM Society*. 2009.

[24] M. P. R. van Moorselaar. “Gas Path Analysis on the GEnx-1B at KLM Engine Services”. MA thesis. Delft University of Technology, 2018.

[25] M. Fokkema. “Transformer-based Outlier Detection in Multivariate Time Series”. MA thesis. Delft University of Technology, 2023.

[26] D. C. Acevedo. “Augmenting Aircraft Engine Flight Data with Generative Adversarial Networks for Fault Detection”. MA thesis. Delft University of Technology, 2024.

[27] M. L. Verbist. “Gas Path Analysis for Enhanced Aero-Engine Condition Monitoring and Maintenance”. PhD thesis. Delft University of Technology, 2017. DOI: 10.4233/UUID:E1079009-84C2-482D-AFE4-E1F9FDE0D137.

[28] Y. Lei et al. “Machinery Health Prognostics: A Systematic Review from Data Acquisition to RUL Prediction”. In: *Mechanical Systems and Signal Processing* 104 (2018), pp. 799–834. DOI: 10.1016/j.ymssp.2017.11.016.

[29] W. P. J. Visser. *Generic Analysis Methods for Gas Turbine Engine Performance: The Development of the Gas Turbine Simulation Program GSP*. S.l.: s.n., 2014. ISBN: 978-94-6259-492-0.

[30] GasTurb GmbH. *GasTurb: Gas Turbine Performance Simulation Software*. Accessed: 2026-04-25. 2024. URL: https://www.gasturb.com/.

[31] NASA Glenn Research Center. *NPSS Numerical Propulsion System Simulation*. NASA Software Catalog ID: LEW-17051-1, Accessed: 2026-04-25. 2026. URL: https://software.nasa.gov/software/LEW-17051-1.

[32] EcosimPro / Empresarios Agrupados Internacional. *PROOSIS: Propulsion Object-Oriented Simulation Software*. Accessed: 2026-04-25. 2024. URL: https://ecosimpro.com/products/proosis/.

[33] S. Ramdin et al. “Systematic Approach for Modelling Modern Turbofan Engines”. In: *Volume 1: Aircraft Engine*. Boston, Massachusetts, USA: American Society of Mechanical Engineers, 2023, V001T01A032. DOI: 10.1115/GT2023-103548.

[34] A. Stamatis et al. “Jet Engine Fault Detection with Discrete Operating Points Gas Path Analysis”. In: *Journal of Propulsion and Power* 7.6 (1991), pp. 1043–1048. DOI: 10.2514/3.23425.

[35] N. Aretakis, K. Mathioudakis, and A. Stamatis. “Non-Linear Engine Component Fault Diagnosis from a Limited Number of Measurements Using a Combinatorial Approach”. In: (2003).

[36] H. Pieters and J. P. Van Buijtenen. “Gas Path Analysis with GSP for the GEM42 Turboshaft Engine”. MA thesis. Delft University of Technology, 2005.

[37] A. J. Volponi and L. Tang. “Improved Engine Health Monitoring Using Full Flight Data and Companion Engine Information”. In: *SAE International Journal of Aerospace* 9.1 (2016), pp. 91–102. DOI: 10.4271/2016-01-2024.

[38] I. M. Sobol'. "Sensitivity Estimates for Nonlinear Mathematical Models". In: *Mathematical Modelling and Computational Experiments* 1.4 (1993), pp. 407–414.

[39] C. Rong, S. OuYang, and H. Sun. "Anomaly Detection in QAR Data Using VAE-LSTM with Multihead Self-Attention Mechanism". In: *Mobile Information Systems* (2022), p. 8378187. DOI: 10.1155/2022/8378187.

[40] Z. Chen et al. "Machine Remaining Useful Life Prediction via an Attention-Based Deep Learning Approach". In: *IEEE Transactions on Industrial Electronics* 68.3 (2021), pp. 2521–2531. DOI: 10.1109/TIE.2020.2972443.

[41] I. Goodfellow et al. "Generative Adversarial Networks". In: *Communications of the ACM* 63.11 (2020), pp. 139–144. DOI: 10.1145/3422622.

[42] C. Little et al. "Generative Adversarial Networks for Synthetic Data Generation: A Comparative Study". In: *arXiv preprint* (2021). arXiv:2112.01925 [cs]. DOI: 10.48550/arXiv.2112.01925.

[43] N. Park et al. "Data Synthesis Based on Generative Adversarial Networks". In: *Proceedings of the VLDB Endowment* 11.10 (2018), pp. 1071–1083. DOI: 10.14778/3231751.3231757.

[44] Y. Raptodimos and I. Lazakis. "Using Artificial Neural Network-Self-Organising Map for Data Clustering of Marine Engine Condition Monitoring Applications". In: *Ships and Offshore Structures* 13.6 (2018), pp. 649–656. DOI: 10.1080/17445302.2018.1443694.

[45] G. Torella. "Expert Systems for Gas Turbine Engines Working Problems". In: *34th AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit*. Cleveland, OH, U.S.A.: American Institute of Aeronautics and Astronautics, 1998. DOI: 10.2514/6.1998-3755.

[46] E. Mohammadi and M. Montazeri-Gh. "A Fuzzy-Based Gas Turbine Fault Detection and Identification System for Full and Part-Load Performance Deterioration". In: *Aerospace Science and Technology* 46 (2015), pp. 82–93. DOI: 10.1016/j.ast.2015.07.002.

[47] W. R. Jacobs et al. "Inter-engine Variation Analysis for Health Monitoring of Aerospace Gas Turbine Engines". In: (2018).

[48] Emoscopes. *Turbofan Operation (lbp). PNG Diagram of a 2-Spool, Low-Bypass Turbofan Engine*. Wikimedia Commons, Accessed: 26 November 2025, CC BY-SA 3.0 / CC BY 2.5 / GFDL. 2005. URL: https://upload.wikimedia.org/wikipedia/commons/8/81/Turbofan_operation_%28lbp%29.png.

[49] W. P. J. Visser et al. *Technical Manual of Gas Turbine Simulation Program*. NLR, 2023.

[50] S. van Vuuren. "Humidity Effects on Turbofan Performance". MA thesis. Delft University of Technology, 2019.

[51] J. Kurzke. "How to Get Component Maps for Aircraft Gas Turbine Performance Calculations". In: *ASME International Gas Turbine and Aeroengine Congress & Exhibition*. American Society of Mechanical Engineers, 1996.

A

Thesis Assignment

Reducing Diagnostic Uncertainty in Gas Path Analysis using Companion Engine performance data

MSc Assignment for Lucas Middendorp
Faculty of Aerospace Engineering, TU Delft, date 23-7-2025

Introduction

KLM Engine Services (ES), a part of the Air France Industries KLM Engineering & Maintenance Group, handles approximately 110 aircraft engines each year. The facility specializes in engines such as the CF6-80E1, CF6-80C2, GEnx-1B, CFM56-7B, and both LEAP-1A and LEAP-1B variants. To support ongoing improvements in diagnostics and maintenance, KLM ES is investigating how existing operational data can be used more effectively within the framework of condition-based maintenance strategies.

Gas Path Analysis (GPA) assesses engine health using sensor data, but limited sensor coverage can cause ambiguity in diagnostics. Existing methods, such as applying degradation trend constraints or incorporating data from auxiliary sensors, can reduce diagnostic uncertainty in specific scenarios. This thesis proposes incorporating fleet-level companion engine data to further improve GPA's effectiveness, especially when sensor coverage is sparse. By using similarities between companion engines and integrating them into GPA, the research aims to constrain diagnostic solutions.

Required skills:

- Understanding of gas turbine performance / aero-engine technology
- Experience in Python, Pandas and Machine Learning/Deep Learning is considered a bonus

Assignment

1. Conduct a comprehensive review of current practices in gas turbine maintenance, including Gas Path Analysis (GPA) and emerging data-driven diagnostic methods.
2. Compare raw sensor data and create (simple) logic identifying sensor malfunctions.
3. Study fleet-wide degradation with a focus on companion engines on shared aircraft.
4. Evaluate whether diagnostic accuracy and reliability can be improved using engine companion data
5. Develop a method to integrate companion engine data analysis in KLM's diagnostics information systems.
6. Provide practical recommendations for implementing the diagnostic method in maintenance workflows, including requirements for data availability and computational tradeoffs.

Optional work package:

7. Apply Multi-Operating Point Analysis (MOPA) on single-engine flight data for GPA.

Duration & Location

The intended duration of this internship thesis is 10 months, starting as soon as possible. The team is based at Schiphol-Oost, with offices located close to the KLM engine shops and hangars.

Report

Results of the work must be reported in English, with a copy of this assignment. The code should be reproducible, usable by other members, and committed to the designated KLM GitHub repository.

Supervision

Professor,
Prof. dr. ir. P. Colonna

Delft University supervisor,
Dr. ir. W.P.J. Visser

Daily supervisors at KLM
Tim Rootliep, Daniel Cisneros Acevedo

KLM Royal Dutch Airlines

B

GSPy Framework

GSPy is a Python-based adaptation of the Gas Turbine Simulation Program (GSP). A detailed understanding of its simulation framework is essential to identify model limitations and enable further extensions. Section section B.1 presents the component structure and implementation, while Section section B.2 describes the simulation approach employed by GSPy.

B.1. Component Structure

This section describes the component-based architecture of GSPy. Each engine component is implemented as a class that inherits from a base component structure. This inheritance hierarchy ensures that all components share a generic `Run()` method, and that child classes can inherit and override this. This object-oriented structure is modular, enabling users to easily customize or add components without altering the overall simulation framework, as illustrated in Figure B.1.

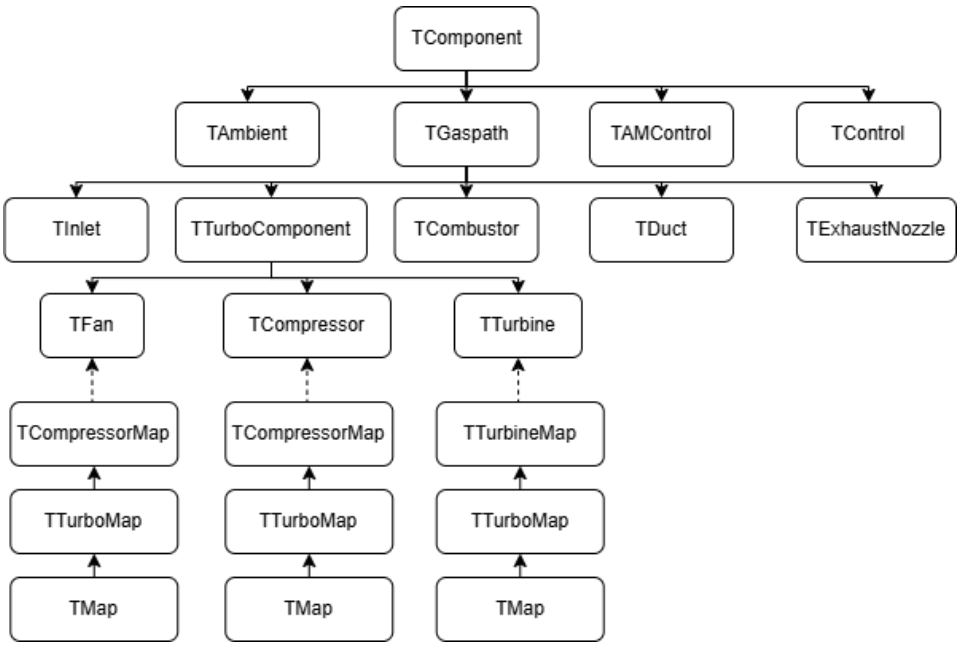

Figure B.1: Class inheritance structure GSPy

B.2. Simulation

This section outlines the procedures used by GSPy to compute the Design-Point (DP) and Off-Design (OD) performance simulations. Together, these steps form the foundation of the GSPy engine perfor-

mance pipeline. The DP simulation defines the reference operating point from which all OD simulations are scaled. Before OD performance can be evaluated, the characteristics of each engine component, such as mass flow rates, pressure ratios, temperatures, and efficiencies, must be established at the reference point. The DP calculation then results in the engine design and represents the cycle reference point (CPR) for OD simulations [29]. Off-Design (OD) simulations evaluate the performance of the design-point model under operating conditions that deviate from the reference state.

In GSPy, a DP simulation is executed using the `Run_DP_simulation()` method. This function initializes the system by assigning the appropriate amount of free states and residual equations, and then iteratively calls the `Run()` method in each component class.

B.2.1. Components

As introduced in the GSPy framework, the engine model is constructed using an object-oriented approach in which each component represents a physical element of the gas turbine. Each component is responsible for computing its thermodynamic state and performance based on incoming flow conditions, while storing the relevant gas path parameters.

The following sections describe the modeling approach and calculation procedures for the components included in the GEnx-1B engine model.

Ambient

Before evaluating the gas turbine performance, the ambient conditions must be established, as they define the inlet boundary state of the engine cycle. Several parameters are provided as input, including altitude, flight Mach number, and a deviation in static temperature relative to standard atmospheric conditions. Static temperature and static pressure may optionally be specified; if either is set to None, the Inlet component computes the missing property using the standard atmosphere model and the prescribed altitude and temperature deviation.

Based on these inputs, the Ambient module constructs a Cantera gas object that contains the total and static thermodynamic properties of the freestream. The flow velocity is subsequently determined from the Mach number and local speed of sound, completing the definition of the inlet state used by the Inlet component.

FuelControl

The Controller is a vital part of the Off-Design calculations, as it regulates the operational parameter that drives the OD schedule. Depending on the selected mode, the controller may prescribe fuel flow, fan speed (N1%), or exhaust gas temperature (EGT). The user provides both the design-point value and an OD sweep definition (start, end, and increment). After the design point `Run()` function of the controller is called, a state and error equation is added for the controller.

Inlet

The Inlet is the first component in the engine model. The `Run()` function of the inlet initializes the `GasIn` and `GasOut` gas states, assigning temperature, pressure, composition and mass flow. The inlet also computes the corrected mass flow via the `GetFlowCorrectionFactor()` function, which non-dimensionalizes the flow based on the ratio of ambient conditions to standard reference conditions. During OD simulations, the Inlet component first updates the corrected mass flow using the inlet mass-flow free state. The mass flow may change under OD conditions and the inlet gas properties are recalculated accordingly.

Additionally, the inlet evaluates the ram drag using the ambient velocity, thereby accounting for the momentum deficit associated with air ingestion. The outlet gas state produced by the Inlet is passed directly to the first compression component.

Fan/LPC

In the GEnx model, the Fan and Low-Pressure Compressor (LPC) are implemented as a single combined component, referred to here as the fan. The fan comprises of two parallel compression paths: the core stream and the bypass (duct) stream. The first step in the `Run()` function is therefore to partition the inlet mass flow according to the design mass flow and prescribed bypass ratio. The flows are given thermodynamic properties using Cantera and the correct mass flow is calculated using the `GetFlowCorrectionFactor()` function.

The `ReadMapAndSetScaling()` is then executed. This calls the `ReadMap()` function to extract performance map data and uses a SciPy grid interpolator to determine the design-point map values of corrected mass flow, pressure ratio, and isentropic efficiency. These values are used to compute the scaling factors applied during OD operation.

Compression is performed using the `Compression()` function from utils.py. Given a pressure ratio, the outlet pressure is computed as $P_{out} = P_{in} \cdot PR$. An isentropic outlet state is then found by imposing constant entropy at the updated pressure. The actual outlet enthalpy is obtained by correcting the isentropic enthalpy rise with the isentropic efficiency:

$$h_{out} = h_{in} + \frac{h_{is,out} - h_{is,in}}{\eta_{is}}$$

The shaft power required by each stream is calculated using the mass flow evaluated at the outlet state. The core and bypass contributions yield the total fan shaft power. Finally, the free states (N_1, β_{core}, β_{duct}, BPR) and residual equations (core corrected mass flow and duct corrected mass flow), before passing `GasOut` to the High-Pressure Compressor.

In OD simulations the `Fan` class first calculates the corrected shaft speed using the `N1 shaft speed` state, and correcting for temperature. The corrected shaft speed and β are used as inputs in the `GetScaledMapPerformance()` function, which then outputs the corrected mass flow, pressure ratio and efficiency. This is done for both the core and duct, and after this the core and duct power are calculated using `Compression()`. Both core and duct power are added to get the total N1 shaft power. The residuals of the core and duct corrected mass flow are then calculated using Equation B.2.

HPC

The HPC, implemented using the Compressor class, begins by computing the compression power using the `Compression()` function. The component then calls `ReadMapAndSetScaling()`, identical to the fan, to load map data and establish scaling factors.

The component subsequently sets its free states and residual equations. If the constant-speed (CS) option is disabled, an additional state representing the shaft rotational speed is introduced. The HPC also incorporates bleed extraction, where bleed mass flow is compressed separately and its impact on shaft power is accounted for explicitly.

In OD simulations the HPC follows the same principle of the fan: First calculating the corrected shaft speed, using that to find the corrected mass flow, pressure ratio and efficiency. The shaft power is calculated using the pressure ratio and efficiency, and finally the mass flow residual error is calculated using Equation B.2.

Combustor

The Combustor class receives several design-point inputs, including exit temperature, pressure ratio, combustion efficiency, fuel inlet temperature, lower heating value (LHV), and chemical composition parameters (HC and OC ratios). Inside the `Run()` function, the fuel composition is first established. The combustor then determines the exit enthalpy and equilibrium product composition using Cantera. These outputs are subsequently used as input to the high-pressure turbine. In OD simulations the fuel flow is enforced using the OD free state W_f rather than the design-point value.

HPT

After the combustor state is defined, the HPT component is evaluated using the `TTurbineMap` class. The spool mechnical efficiency is defined, and a root-finding procedure (SciPy `optimize.root` function is used to find the pressure ratio that yields the shaft power required for the HPC. This is done using the `pressure_ratio_for_turbine()` function, which calculates the power the turbine produces using the `TurbineExpansion()` function. Once an appropriate pressure ratio is found, the map scaling parameters are computed via `ReadMapAndSetScaling`.

The HPT requires one state (β) and two residual equations:

1. The corrected inlet mass flow must match the design-point corrected mass flow from the turbine map.
2. The turbine shaft power must equal the power required by the HPC.

The HPT OD calculations are similar to that of the compressors. First the corrected shaft speed is calculated using the free state, this is used to get the corrected mass flow, pressure ratio and efficiency. The shaft power generated by the turbine is calculated using `TurbineExpansion()`. Finally, the two residuals for corrected mass flow and shaft power are calculated. The residual for corrected mass flow is calculated by subtracting the shaft power required for the HPC (`shaft.PW_sum`) by the shaft power produced by the HPT. This is then divided by the DP PW to get the residual.

LPT

The Low-Pressure Turbine is also implemented using the `TTurbine` class and follows the same computational steps as the HPT. Its residual equations ensure that the LPT supplies the shaft power demanded by the fan, and that the corrected inlet mass flow matches the corrected mass flow in the LPT map.

The LPT follows the same OD sequence as the HPT. Its residuals enforce the corrected mass-flow agreement and the N1 spool power balance.

Exhaust Nozzle

The GEnx-1B employs two exhaust nozzles, one for the bypass flow and one for the core flow. Each nozzle is implemented as a separate instance of the `TExhaustNozzle` class. The nozzle first calculates the exit velocity through `calculate_exit_velocity()`, from which the corresponding Mach number is obtained.

If the computed Mach number exceeds unity, a root-solver is invoked to determine the exit pressure that produces a choked flow condition (Mach = 1). After determining the exit pressure and temperature, the component formulates its residual equation enforcing conservation of mass through the nozzle, which will be used in Off-Design simulations.

Both nozzles follow a similar OD procedure. Because nozzle areas remain constant, the design-point throat area is used to compute the expansion through `calculate_expansion_to_A()`. This routine determines the throat pressure, temperature, mass flow, and velocity. The nozzle mass-flow residual is then enforced, and the thrust is evaluated using the throat mass flow, exit velocity, throat area, and exit pressure.

B.2.2. Off-Design Controller

Off-Design (OD) simulations evaluate the performance of the design-point model under operating conditions that deviate from the reference state. These deviations may include changes in altitude, static pressure, ambient temperature, fuel flow, N1 spool speed, or exhaust gas temperature (EGT). To ensure that the prescribed OD variable is satisfied, a controller is introduced. Conceptually, the controller adds an additional residual equation (or constraint) that must be satisfied by the nonlinear solver, in addition to the nine residuals already defined in the model.

For example, when the FuelControl component is configured to regulate the N1 spool speed, the controller imposes the following constraint:

$$r_{control} = \frac{N1\%_{demand} - N1\%_{value}}{N1\%_{DP}} = 0 \tag{B.1}$$

Where the $N1\%_{demand}$ is the value set by the Control class, the $N1\%_{value}$ is the calculate value after an OD simulation is run, and $N1\%_{DP}$ is the design-point value of $N1\%$.

Furthermore, a free state must be defined to ensure the number of free states equals the number of error equations. For the FuelControl Controller the fuel flow W_f is set as free variable. The complete set of states and error equations for the GEnx-1B closed-loop off-design simulation is summarized in Table B.1.

In the off-design model, the components (Fan, HPC, LPT, HPT) must satisfy both flow and work/shaft balance. The residual equations listed in Table B.1 enforce these balances. However, to evaluate mass flow, pressure ratio and efficiency produced by each component at off-design conditions, the solver must rely on component maps.

Component	State	Description	Residual equations
Inlet	$W_{c,inlet}$	Corrected inlet mass flow	-
FuelControl	-	Fuel control schedule	r_{control}
Fan	N_1, β_{core}, β_{duct}, BPR	N1 spool speed, core map coordinate, duct map coordinate, bypass ratio	$r_{\text{core,flow}}$, $r_{\text{duct,flow}}$
HPC	N_2, β_{HPC}	N2 spool speed, HPC map coordinate	$r_{\text{HPC,flow}}$
Combustor	W_f	Fuel flow rate	-
HPT	β_{HPT}	HPT map coordinate	$r_{\text{HPT,flow}}$, r_{shaft,N_2}
LPT	β_{LPT}	LPT map coordinate	$r_{\text{LPT,flow}}$, r_{shaft,N_1}
Exhaust Duct	-	Duct nozzle mass flow	$r_{\text{exh,flow}_{\text{duct}}}$
Exhaust Core	-	Core nozzle mass flow	$r_{\text{exh,flow}_{\text{core}}}$

Table B.1: Overview of free state variables and residual equations for the GEnx-1B engine

B.2.3. Component Maps

Component maps describe the performance characteristics of turbomachinery elements in a non-dimensional form. Each map is expressed as a relationship between:

- Corrected shaft speed (N_c)
- Corrected mass flow (W_c)
- Pressure ratio (PR)
- Efficiency (η)
- Map Coordinate (β)

The map coordinate β is introduced to mitigate interpolation errors that arise in regions where corrected speed lines become nearly horizontal or vertical [51]. In GSPy, the corrected mass flow, pressure ratio and efficiency are found using the corrected shaft speed and map coordinate β. The values are stored in a .map file, which consists of a three sub-tables (corrected mass flow, pressure ratio, efficiency) and β and N_c as the X and Y coordinates. Then a continuous representation of the map is constructed using the scipy `RegularGridInterpolator` function.

This procedure is applied to the maps of the fan core, fan duct, HPC, HPT and LPT. In the turbine components the procedure does not require the pressure ratio to be found. The residual equations for the mass flow are thus the following:

$$r_{\text{flow}} = \frac{W_{\text{c,out}} - W_{\text{c,in}}}{W_{\text{c,des}}} = \frac{f(\beta, N_c) - W_{\text{c,in}}}{W_{\text{c,des}}} = 0 \tag{B.2}$$

Where f(β, N_c) represents the interpolated map mass flow. The shaft speed residual equations are the following, where the shaft power of the compressor must equal the shaft power delivered by the turbine, taking in account the mechanical efficiency of the shaft.

$$r_{\text{shaft}} = \frac{-P_{\text{comp}} + \eta_{\text{mech}} P_{\text{turb}}}{P_{\text{des}}} = 0 \tag{B.3}$$

The system of free states and residual equations now form a closed set of equations. In the GEnx-1B model, the system consists of 10 unknowns, and 10 equations to solve. GSPy employs the Newton-Raphson method to iteratively solve the system of equations, containing the free state vector **s** and the residual vector **R(s)**.

$$\underbrace{\begin{bmatrix} r_{\text{core,flow}} \\ r_{\text{duct,flow}} \\ r_{\text{HPC,flow}} \\ r_{\text{control}} \\ r_{\text{HPT,flow}} \\ r_{\text{shaft},N_2} \\ r_{\text{LPT,flow}} \\ r_{\text{shaft},N_1} \\ r_{\text{exh,flow}_{\text{duct}}} \\ r_{\text{exh,flow}_{\text{core}}} \end{bmatrix}}_{R(\mathbf{s})} = \mathbf{0} \quad \Longleftrightarrow \quad \underbrace{\begin{bmatrix} W_c \\ N_1 \\ \beta_{\text{core}} \\ \beta_{\text{duct}} \\ \text{BPR} \\ N_2 \\ \beta_{\text{HPC}} \\ W_f \\ \beta_{\text{HPT}} \\ \beta_{\text{LPT}} \end{bmatrix}}_{\mathbf{s}}$$

B.2.4. Adaptive Modeling

In addition to standard OD analysis, GSPy incorporates Adaptive Modeling (AM), a feature inspired by the implementation in GSP. AM enables estimation of component health parameters based on measurement data (e.g., test-cell, on-wing, or simulated measurements). Using these measurements, GSPy determines the deviation of component parameters (efficiency, corrected mass flow, pressure ratio) from their nominal design-point values.

Adaptive Modeling is implemented through the `TAMcontrol` class. During the design-point execution, `TAMcontrol` adds additional free states and residual equations, depending on the number and type of measurements provided. Each free state represents a map-scaling factor, either efficiency, corrected mass-flow or pressure ratio. After obtaining the corrected mass flow, efficiency, and pressure ratio from the nominal characteristic map, the model applies the adaptive map-scaling factors within the `GetScaledMapPerformance()` function. These adaptive factors are multiplied with the existing baseline scaling parameters, resulting in modified corrected mass flow, efficiency or pressure ratio that reflect the deteriorated state of the component. Without AM, these deterioration scaling factors still exist, but are set to one.

The added residual equations due to the AM class are calculated using the following equation:

$$r_i = \frac{\text{Output}_i - \text{Measured}_i}{\text{Model}_{\text{des,i}}} \tag{B.4}$$

Where i is the parameter selected and $Output_i$ is the corresponding GSPy OD simulation output. An example is given below of the AM residual vector and the AM state vector of the GEnx-1B. Since the on-wing measurement set of the GEnx-1B only contains 6 measurements, the amount of map modifiers must also be 6.

$$\underbrace{\begin{bmatrix} r_{\text{T25}} \\ r_{\text{P3}} \\ r_{\text{T3}} \\ r_{\text{T49}} \\ r_{\text{N}_2} \\ r_{\text{W}_f} \end{bmatrix}}_{R(\mathbf{s})} = \mathbf{0} \quad \Longleftrightarrow \quad \underbrace{\begin{bmatrix} \text{Fan SF } \eta \\ \text{Fan SF } W_c \\ \text{HPC SF } \eta \\ \text{HPC SF } W_c \\ \text{HPT SF } \eta \\ \text{HPT SF } W_c \end{bmatrix}}_{\mathbf{s}}$$

These AM states and residual equations are appended to the ten free states and ten residual equations already present in the baseline GEnx-1B model, thereby forming a fully closed system for adaptive parameter estimation. Once the augmented system is defined, a standard DP and OD simulation can be executed. The Newton–Raphson solver then iteratively adjusts the six map-scaling factors such that the simulated outputs match the on-wing measurements as closely as possible, completing the adaptive modeling process.

C

Non-Converging Subset Analysis

Some subsets converge while other do not. An analysis into the numerical equations reveals this behavior is not random, but comes from the thermodynamic coupling between the components. Below, explanations are given for the non-convergence of all subsets. The non-converging subsets are subsets 13-28, and are illustrated below in Figure C.1.

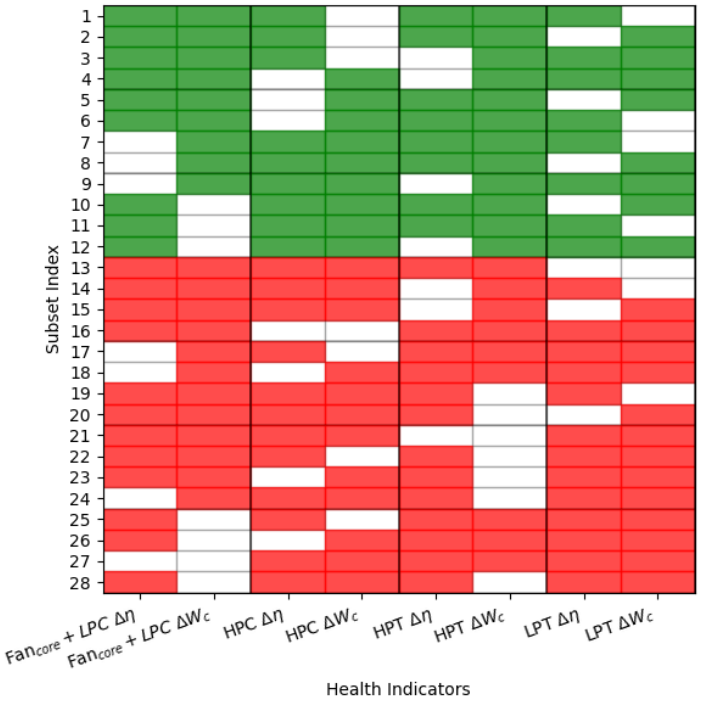

Figure C.1: Subsets of health indicators

C.1. HPT Flow Capacity

In all working subsets, the HPT flow capacity deviation is included. Subsets without the HPT flow capacity (subsets 19-24 & 28) do not converge. This stems from the core energy balance. In the core section of the engine, the air mass flow W_{air} is constant. The energy equation for the combustion chamber is the following:

$$W_f \cdot LHV \cdot \eta_{cc} = W_{air} c_p (T_4 - T_3) \tag{C.1}$$

Since the fuel flow W_f and HPC outlet temperature T_3 are measured, the turbine inlet temperature T_4 only depends on the air mass flow W_{air}. In addition, the corrected air mass flow of the HPT W_{c4} is defined as:

$$W_{c4} = W_{air}\sqrt{\frac{T_4}{T_{amb}}}\frac{P_{amb}}{P_4} \tag{C.2}$$

Here, the turbine inlet pressure P_4 is fixed since P_3 is measured, and the pressure loss over the combustor is assumed constant. If W_{c4} is not allowed to vary, the solver must satisfy both:

- The combustion energy equation (Equation C.1) which determines T_4
- The turbine corrected flow equation (Equation C.2) which depends on T_4

This creates a constraint in the form of:

$$W_{air} \cdot T_4 = W_{air} \cdot \sqrt{T_4}$$

Which cannot be satisfied under deterioration. This causes the system to become inconsistent and fails to converge. For this, the HPT flow capacity delta must always be included in the subsets [36].

C.2. Remaining Non-Converging Subsets

The remaining non-converging subsets show a clear pattern. If the engine is divided into a compressor side (fan core + LPC and HPC) and a turbine side (HPT and LPT), at least three out of the four corresponding health indicators on each side must be included. Subsets that violate this condition consistently fail to converge.

This behavior can be explained by considering the balance between the number of unknown states and the constraints imposed by the measurements and residual equations. Each engine component is governed by a set of coupled equations linking pressures, temperatures, spool speeds, and mass flows. These relations constrain combinations of health indicators rather than allowing them to vary independently.

When both health indicators of a component are excluded, the model loses the flexibility required to match key measurements such as temperatures, pressures, and fuel flow. As a result, the remaining health indicators must compensate for multiple physical effects simultaneously, which introduces conflicting requirements. This leads to an overconstrained system in which the residual equations cannot be satisfied. More generally, excluding too many health indicators from one side of the engine reduces the available degrees of freedom below what is required to satisfy the coupled thermodynamic relations. This results in a non-converging system, where deterioration of too many health indicators prevent convergence of the solver.

In conclusion, the requirement to include at least three out of four health indicators per engine side is a direct consequence of the coupled nature of the thermodynamic equations and the need to maintain sufficient degrees of freedom. Subsets that do not satisfy this condition lack the flexibility required match the measurements with their health indicator set, and therefore do not converge.

D

Companion Engine Analysis Validation

A fundamental assumption in companion engine analysis is that both engines operate under comparable conditions. This assumption is essential, as deviations in health indicators between engines should primarily reflect component deterioration rather than differences in operating conditions. Due to the data extraction method used, the data is first filtered per point, before being separated into a left and right engine dataset.

Figure D.1 illustrates the operating conditions of the main and companion engines. The Mach number exhibits strong similarity between both engines, with limited scatter across the operating range. Similarly, the corrected fan speed (N1c) follows a consistent trend, although a limited number of outliers deviate from the diagonal relationship.

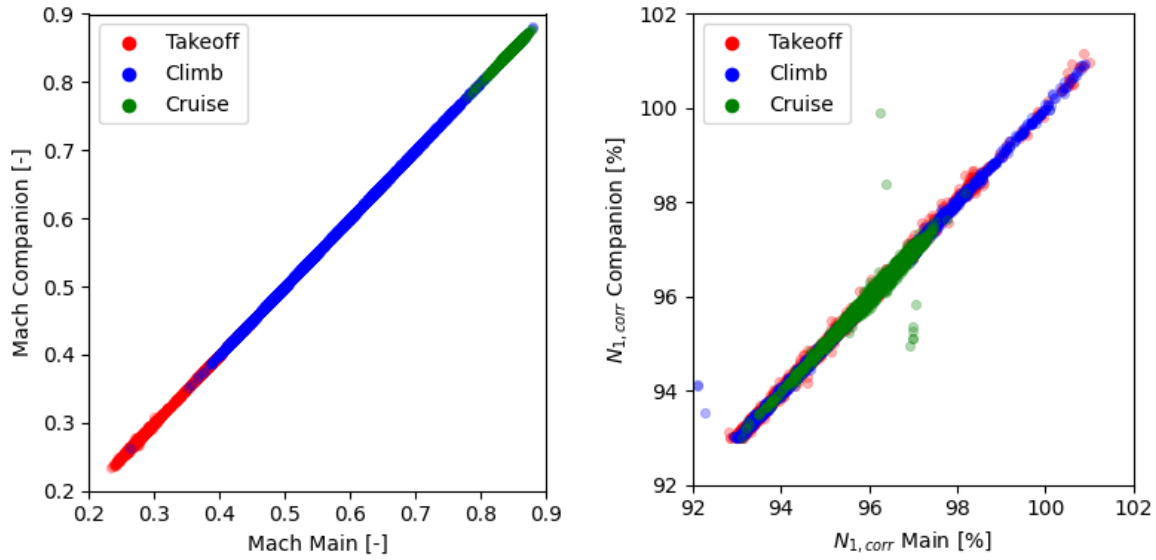

Figure D.1: Main and companion engine operating conditions

An additional assumption underlying companion engine analysis is that both engines exhibit comparable deterioration behavior over time. Significant deviations from this assumption may indicate abnormal degradation in one of the engines. Consequently, sensor measurements, which serve as indirect indicators of engine health, are also expected to remain consistent between the two engines.

Figure D.2 presents the filtered sensor measurements for both the main and companion engines. The

majority of measurements closely align with the diagonal, indicating a strong level of similarity between the two engines. This suggests that both engines are subjected to similar operating conditions and exhibit comparable levels of deterioration. Therefore, this validation confirms the use of the companion engine assumption.

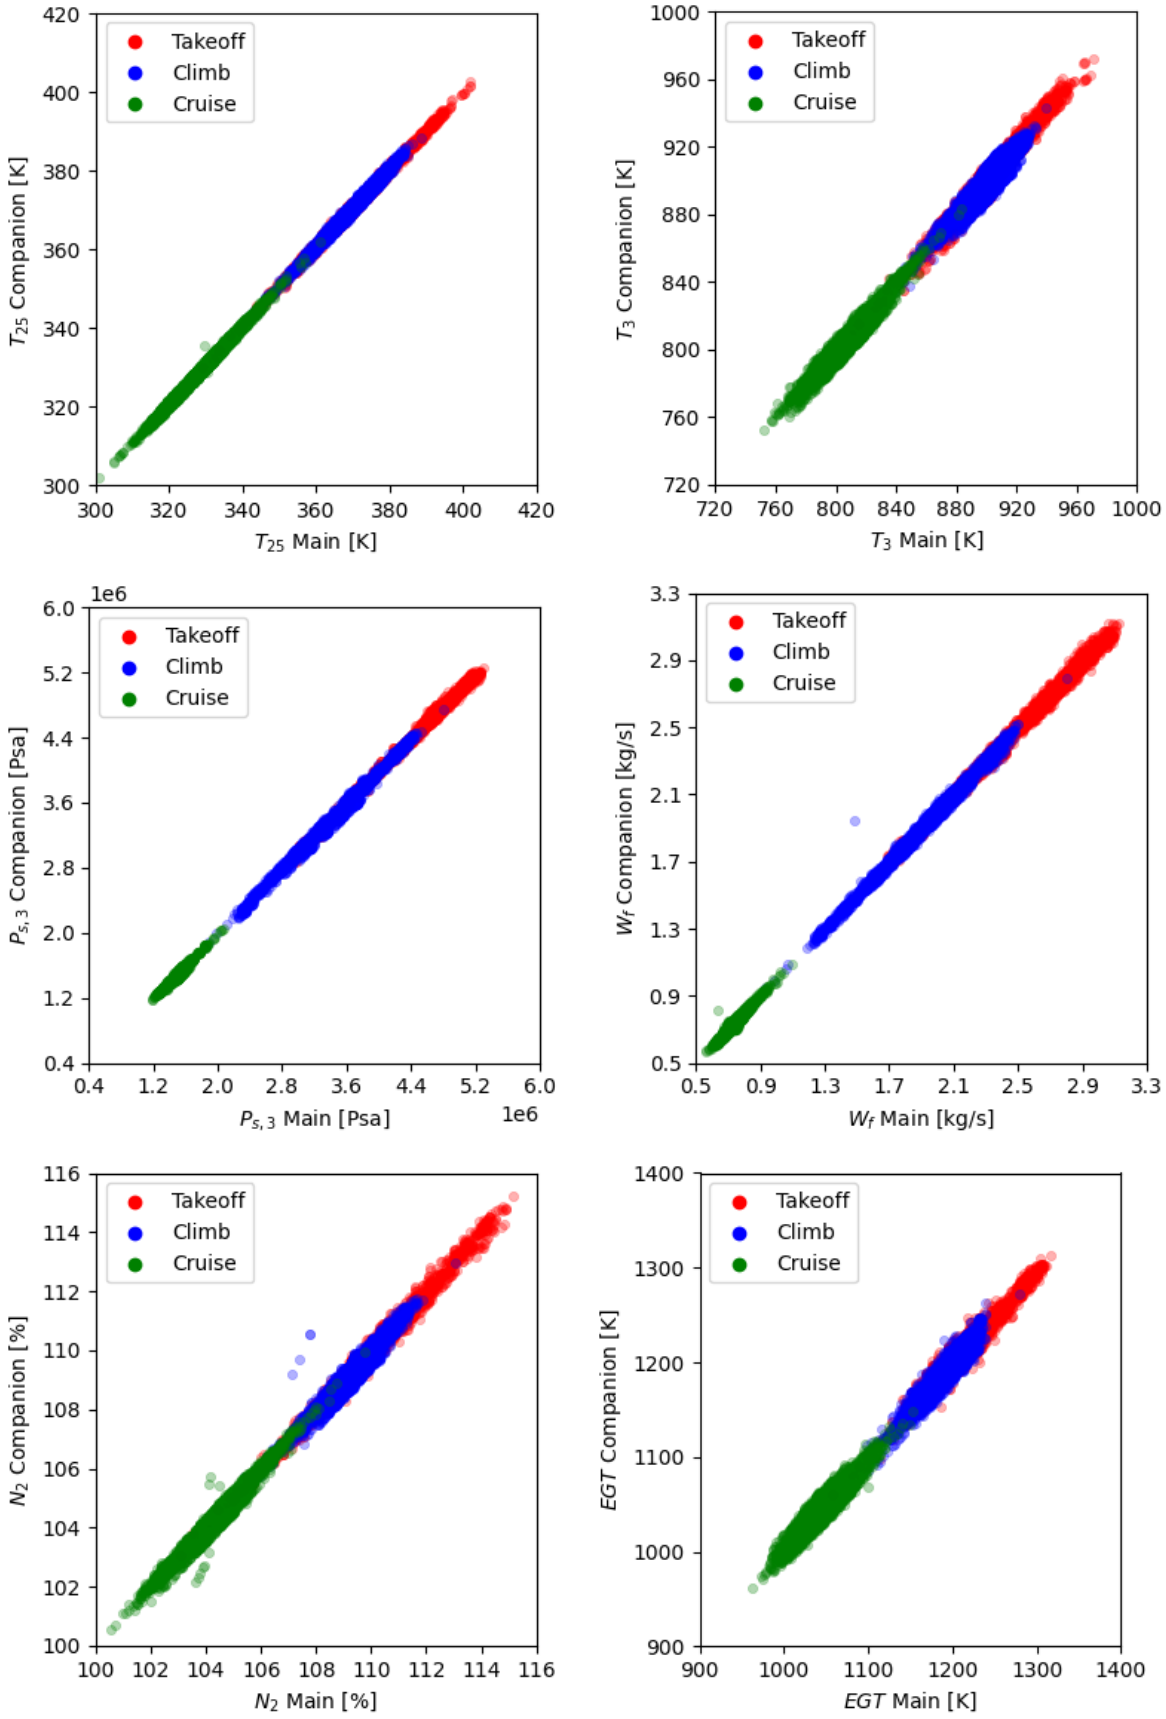

Figure D.2: Main and companion engine sensor measurements

E

Subset Results Validation Case 1

The adaptive modeling results of the subsets regarding the first validation case can be seen below.

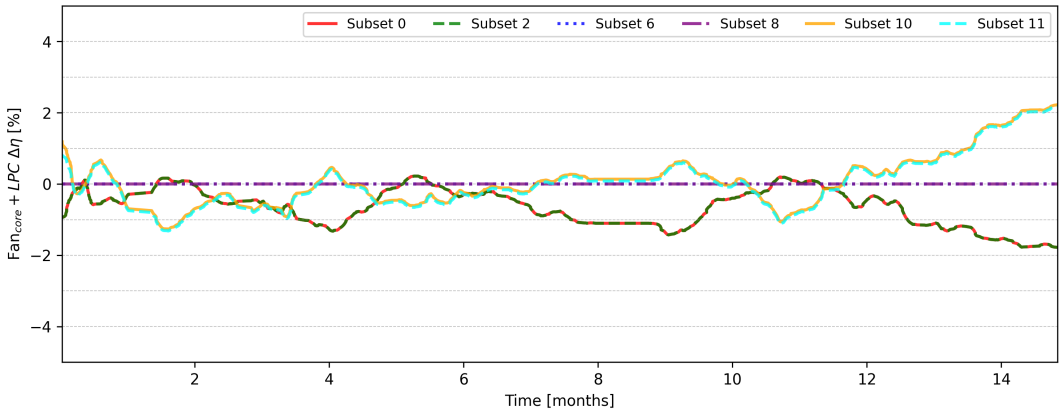

Figure E.1: Fan core + LPC efficiency subset results rolling average

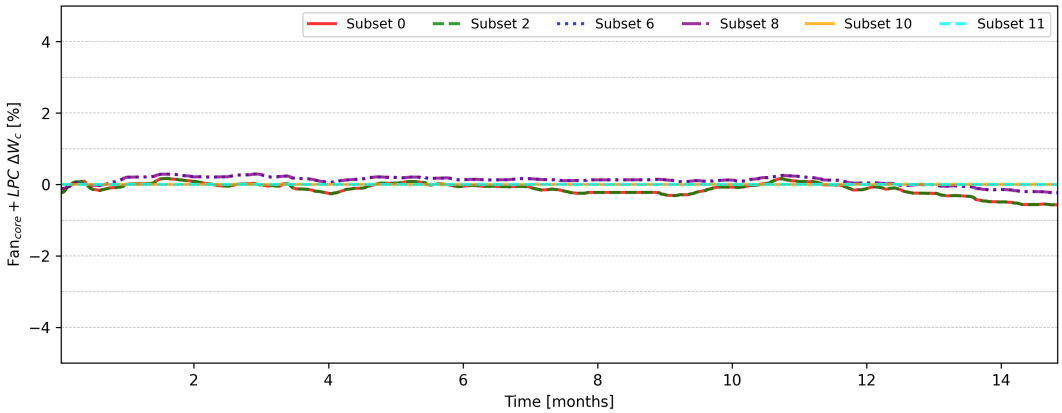

Figure E.2: Fan core + LPC flow capacity subset results rolling average

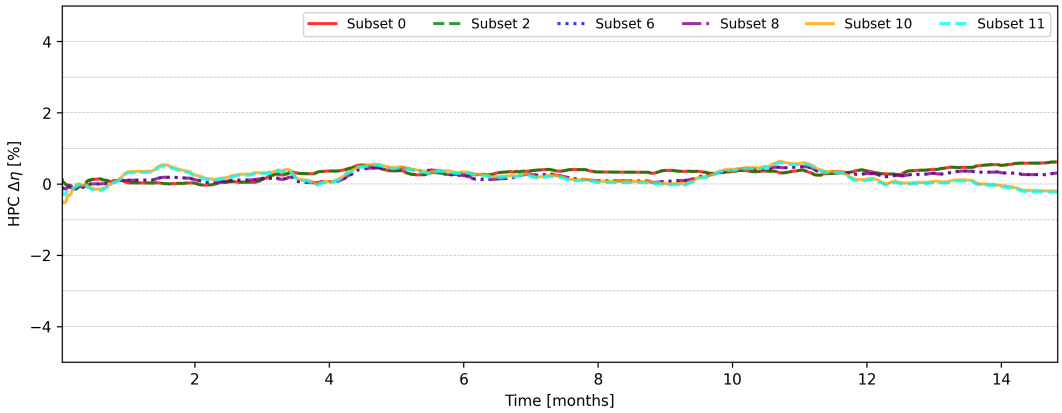

Figure E.3: HPC efficiency subset results rolling average

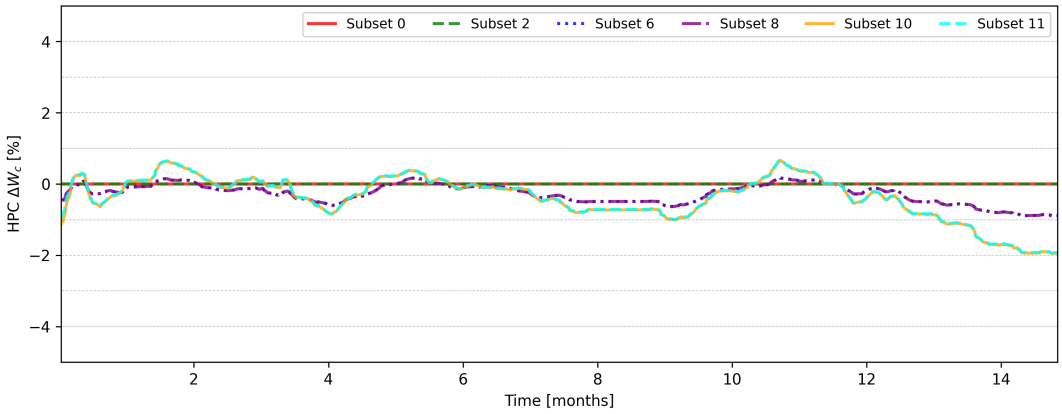

Figure E.4: HPC flow capacity subset results rolling average

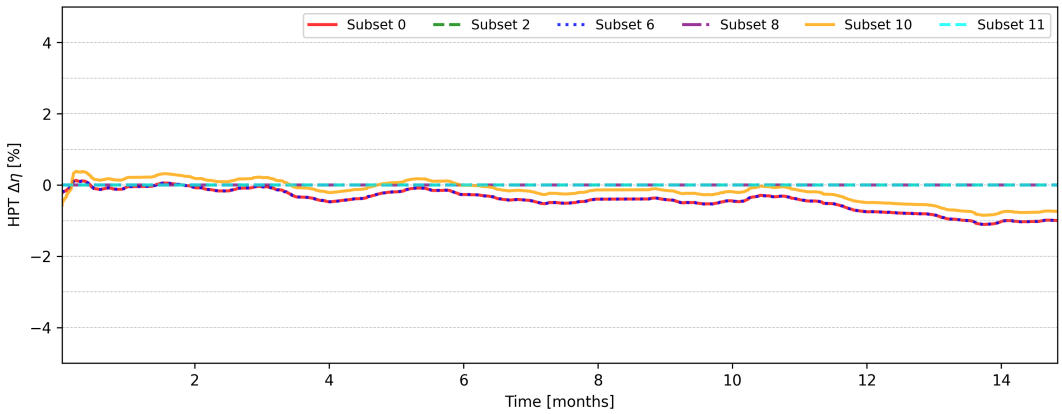

Figure E.5: HPT efficiency subset results rolling average

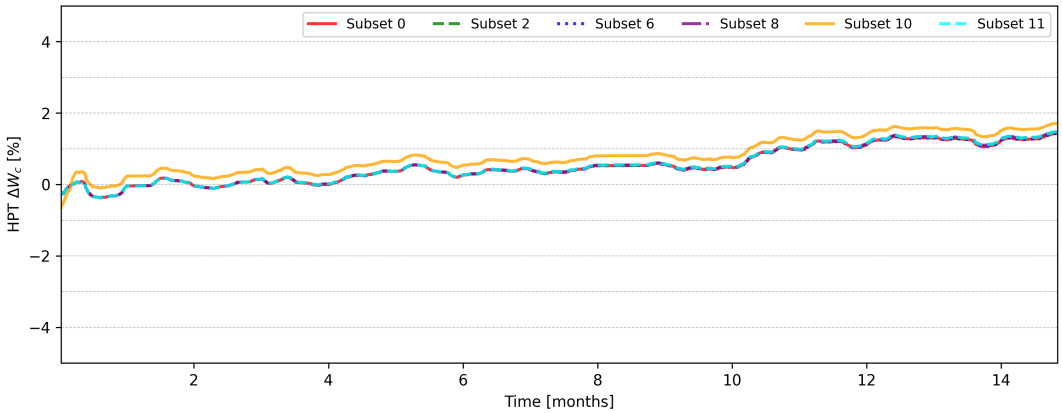

Figure E.6: HPT flow capacity subset results rolling average

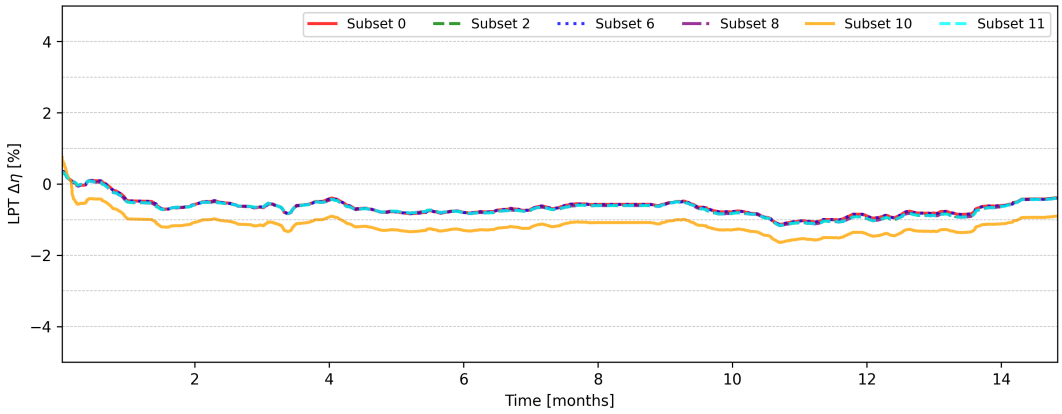

Figure E.7: LPT efficiency subset results rolling average

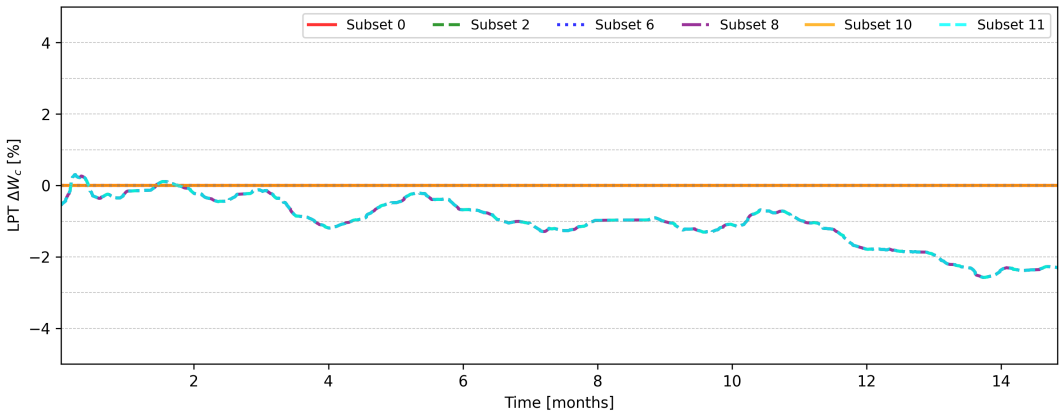

Figure E.8: LPT flow capacity subset results rolling average