Doctoral Dissertation
Doctoral Program in Mechanical Engineering (38^{th}cycle)

An Industrial PHM Framework Aimed at Condition-Based Maintenance and Fleet Management of Advanced Jet Trainers

By

Leonardo Baldo

******

Supervisor(s):
Prof. Massimo Sorli, Supervisor
Prof. Andrea De Martin, Co-Supervisor
Dr. Mathieu Terner, Industrial Tutor

Doctoral Examination Committee:
Prof. Francesco Braghin, Referee, Politecnico di Milano
Prof. Ian Jennions, Referee, Cranfield University

Politecnico di Torino
2026

Declaration

I hereby declare that, the contents and organization of this dissertation constitute my own original work and does not compromise in any way the rights of third parties, including those relating to the security of personal data.

Leonardo Baldo
2026

* This dissertation is presented in partial fulfillment of the requirements for **Ph.D. degree** in the Graduate School of Politecnico di Torino (ScuDo).

The doctoral thesis was written at the end of the doctoral programme funded by PNRR, Missione 4, componente 2 “Dalla Ricerca all’Impresa” – Investimento 3.3 “Introduzione di dottorati innovativi che rispondono ai fabbisogni di innovazione delle imprese e promuovono l’assunzione dei ricercatori dalle imprese”, tramite il Decreto Ministeriale n. 352 del 9 aprile 2022.

Finanziato dall'Unione europea NextGenerationEU

Ministero dell'Università e della Ricerca

Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA

A mio zio Franco e a mia nonna Rosella, che mi hanno trasmesso il valore dell'apprendimento e alimentato la mia curiosità, molto prima che sapessi cosa fosse un dottorato.

Acknowledgments

A PhD teaches you a particular relationship with time: long stretches of uncertainty, sudden moments of clarity, and the discipline of returning to the same problem until it clicks. This thesis was shaped by people who helped at different points along that timeline, in ways that mattered. I am grateful to all of them.

I would first like to thank my supervisor, Prof. Massimo Sorli, his trust in my decisions and direction allowed this work to develop around the questions I felt were worth pursuing. I am deeply grateful to my co-supervisor, Prof. Andrea De Martin. Without his constant support, I would not have made it through this journey. His availability, insight, and the many discussions we had over the years were essential in refining ideas into results. I also thank my industrial tutor, Dr. Mathieu Terner, for his strong involvement throughout the project and for consistently pushing for precision and clarity. His feedback was consistently thorough, and I have come to recognize, especially in hindsight, its full value. I have internalized that rigor, and it has shaped the standards I now apply to my work and to the way I communicate it.

I am grateful to Prof. Paolo Maggiore and Prof. Matteo Dalla Vedova, my advisors during my Bachelor's and Master's theses. Working with them introduced me to PHM and shaped my early research interests. It has been a pleasure to collaborate over the years, and to share many conversations beyond the office, often over dinner. I look forward to continuing our collaboration in the years ahead.

I also thank my friends Gaetano and Ottavio. Working in related fields, we have shared countless discussions where actuators somehow become the main topic every time. I truly value their friendship and the exchange of ideas that comes with it. I would also like to thank Roberto and Matteo. Even as our careers evolved, I have appreciated their support and the many conversations we shared. I hope we will all reunite in Mola. My thanks also go to Alessandro, a friend I truly value, and someone I have always been able to count on.

My visiting PhD period in Santiago, Chile was one of the most formative parts of my PhD. I owe special thanks to Prof. Marcos Orchard, whose mentorship shaped my work in a way that is difficult to describe. He believed in my ideas and gave me the space to develop them, strengthened my vision, and helped me make choices that have had a lasting effect on my career direction. I am also grateful for his friendship, and I truly hope this is only the beginning of a long collaboration. In Chile, I also had the pleasure of meeting Jorge, Ricardo, and Bruno, who became not only valued colleagues but also friends. I sincerely hope there will be many more opportunities para tomar un matecito y jugar un truco juntos.

I also thank my longtime friends Fra, Ester, and Andrea for their friendship and for being a steady point of reference throughout these years. To my parents and Sofia, I am deeply grateful for your unwavering support, for always making space to listen, and for your constant belief in me. I am also grateful to all those I may have unintentionally omitted: your help, in ways both large and small, has been part of this journey.

Executive Summary

The present industrial PhD project addresses a strategic challenge for the aviation sector: how to transform large volumes of operational aircraft data into actionable Prognostics and Health Management (PHM) capabilities for safety-critical flight control actuators, starting from legacy platforms. Developed in partnership between Politecnico di Torino and Leonardo S.p.A. and co-funded under the Italian Piano Nazionale di Ripresa e Resilienza (PNRR) within the EU's NextGenerationEU framework, the work aligns with national and European priorities on digitalisation, innovation and safety-critical infrastructure.

From a strategic perspective, the project responds to three converging drivers. First, aviation is under pressure to deliver higher availability and lower life-cycle costs while maintaining uncompromised safety and regulatory compliance. Second, the proliferation of on-board Health and Usage Monitoring Systems (HUMS) and ground information systems generates data that remain largely under-exploited for prognostics and decision support. Third, the state of the art in AI and machine learning for PHM shows a rapidly expanding body of methods, but only a limited fraction of truly deployed, auditable solutions for real industrial fleets. Against this backdrop, the thesis positions PHM as an enabling capability that must integrate with existing maintenance processes, standards, and operations.

The technical core of the project is the design and implementation of a PHM framework for the electro-hydraulic actuators of an Advanced Jet Trainer (AJT) flight control system, using exclusively operational data from a fleet in service. The work demonstrates that a legacy fleet can be equipped with a structured PHM capability by systematically exploiting HUMS files, equipment registers, maintenance logs and associated documentation. The framework is organised into three functional layers: (i) data management, Extraction–Transformation–Loading (ETL) and pre-processing; (ii) data processing, feature engineering and probabilistic health prediction; and (iii) health management and maintenance decision support. This layered architecture is deliberately technology-agnostic and compliant with existing HUMS infrastructure, making it transferable to other components and case studies.

A key outcome is the scalable ETL workflow that harmonises heterogeneous sources into a consistent data model at flight and component level. This addresses a recurring bottleneck identified in the literature review: most AI/ML-based PHM studies assume clean, curated datasets that do not reflect industrial realities. By

contrast, the proposed workflow deals explicitly with operational data. On top of this foundation, the project develops degradation-sensitive features and flight-level indicators that link operational usage to actuator stress, including mission profile clustering to provide operation-aware future simulation. These features feed probabilistic health prediction models that provide risk trends, while explicitly accounting for future mission profiles and fleet heterogeneity.

Crucially, the framework envisions and facilitates the process of transforming health information into informed decisions. In the last chapter, the thesis outlines and analyses maintenance strategies and fleet management outcomes, both short–mid term (e.g., improved scheduling, prioritisation of removals, enhanced troubleshooting) and long term (e.g., fleet renewal decisions, tail number suggestions). In summary, this thesis delivers a concrete, industry-grade PHM framework for legacy flight control actuators and a strategic understanding of AI/ML trends in prognostics.

Abstract

The integration of equipment operational data into decision-making processes is transforming asset management, laying the foundation for smarter and optimized operations. Prognostics and Health Management (PHM) strategies combine operational monitoring and historical data within a unified framework to assess system behavior and predict future health status leveraging advanced data processing techniques. This approach benefits Original Equipment Manufacturers (OEMs), operators, and maintenance technicians by enabling tailored maintenance planning and increased fleet status awareness. In particular, the implementation of optimized scheduling strategies for maintenance operations facilitates improved asset availability and operational readiness, thereby reducing total operational expenditures and increasing system performance.

The possibility of the integration and development of PHM approaches on historical databases or datasets derived from already operational equipment is one of PHM many selling points. In fact, in this case, the operation of legacy equipment can be enhanced without the need to incorporate new sensors, utilizing the signals already monitored and thus eliminating the need of long and expensive certification phases. The main challenge when dealing with operational equipment is to envision a PHM system able to infer, extract, and interpret operational multimodal data which was not originally and systematically recorded for PHM purposes but that was instead captured for a diverse array of applications, such as Structural Health Monitoring (SHM), control loops, diagnostic purposes, etc. It is then essential to approach the analysis of these sources from an completely different point of view, trying to extract prognostic value from an existing set of signals to analyze prognosable failures. In fact, when a system is prognosable, PHM strategies can be built on existing monitored signals, enabling new capabilities without equipment modification.

However, the effective utilization of PHM strategies on operational industrial data faces considerable challenges, especially in terms of data quality and representativeness. Often, the needed signals are not monitored, sampling frequencies are insufficient, and data sources are inconsistent or isolated in different siloed databases. Due to the limited coverage of these areas in existing literature, this thesis encounters several challenges stemming from the innovative nature of the questions posed and the lack of comparable studies available. Within this high-risk, high-reward context, this thesis develops a PHM framework utilizing operational, historical, and logistic data from a fleet of Advanced Jet Trainers (AJTs). The study focuses on the Horizontal Tail (HT) Electro-Hydraulic flight control Actuator (EHA) of the AJT, which underwent comprehensive analysis to collect and analyze all relevant data sources.

The research methodology involved an extensive data collection campaign, culminating in an Extraction-Transformation-Load (ETL) pipeline, followed by detailed Exploratory Data Analysis (EDA) to identify data-specific strategies. Algorithms were developed and tested to identify influential features created during feature engineering. In this context, flight clustering was performed to enable mission-informed trend propagation in the future. The features were projected into the future using Monte Carlo simulations and compared against a dynamic probabilistic hazard distribution. This process generated a Probability of Failure (PoF) distribution over time, which was integrated to create a risk trend. This risk trend ultimately drove the maintenance optimization framework, divided between short-term and long-term predictions within an explainable pipeline.

In summary, this thesis positions itself at the interface between industrial challenges and academic advances in PHM for EHAs, with the aim of bridging the gap between deployable solutions and state-of-the-art research. Beyond the methodological contributions, it highlights how a structured PHM framework can provide the foundation for concrete operational benefits, from improved situational awareness to more informed maintenance planning. The concepts and tools developed here provide a foundation for future developments that will further exploit operational outputs, enabling their progressive integration into routine maintenance workflows and flight operations, and their adaptation to different customers, fleets, and mission profiles.

Preface

During my three years of PhD research, I immersed myself in the fields of Prognostics and Health Management (PHM) and Condition-Based Maintenance (CBM) – those very topics that have fueled my professional trajectory since my Bachelor's thesis. I firmly believe in the power of appropriate system and equipment monitoring to improve operator awareness, leading to positive outcomes in terms of performance, cost-effectiveness, and safety. Identifying more as a "maintenance technician" than a designer, I have always been drawn to the practical aspects of making things work, repairing them, and optimizing, rather than just conceiving them. Therefore, I viewed this PhD as a perfect opportunity to merge my professional background with my passion: to deepen my knowledge while maintaining a strong connection to aviation, contributing to its safety and efficiency through practical, implementable solutions. It was.

I think my growth over the past three years can be summarized by a few phrases, which have become the pillars through which I approach problems to find solutions.

The first mantra is "Do the best with what you have". Rarely does a project come equipped with every component exactly as ideally needed. In such situations, success depends on making the most of the resources at hand. It is about adapting creatively, leveraging what is available, and turning limitations into opportunities. Analyzing operational fleet data can reveal both its advantages and challenges. Many researchers and data scientists in the aviation field aspire to work with real-life data, drawn by its inherent fascination and the opportunity to examine actual flights conducted by professional pilots in genuine aircraft. However, the challenges associated with these data are significant, as crucial signals for analysis are often not recorded or are logged at insufficient frequencies. This research project aimed to explore the extent to which insights could be derived from sensors and information related to a production aircraft without adding new sensors.

The second mantra of my PhD project has been "Less is more". I have always felt represented by this phrase and I have completely embraced it in my PhD journey by keeping things simple and adding complexity only when needed. This idea is represented in the KISS approach: Keep it simple, stupid!, first adopted by the United States Navy and then formalized as a coding best practice. In a world where AI is used everywhere also when it is not needed, I believe it is a tool and not a solution. One may be surprised that the core of the presented framework is not developed with complex Deep Learning (DL) pipelines. We chose to stick to the ground and provide a solution without the need of adopting complex, non-explainable, hardly transferrable architecture in spite of the "As simple as possible, as complex as necessary" principle. Again, less is indeed more.

The third phrase is "Done is better than perfect because perfect never gets done". As an engineer, I consistently strive for the best possible solution. However, I have come to understand that perfection is often unattainable. At some point, it is necessary to establish a boundary, finalize what has been accomplished, and move forward. Although the solution at hand may not be the ideal one, it is the result we have achieved, and we must place our trust in it and proceed.

This journey provided me with a hybrid set of competences. In fact, in PHM projects, it is common to delineate the roles of the Subject Matter Expert (SME) and the Data Scientist. The SME typically holds expertise in the system and process industrial knowledge, while the data scientist focuses primarily on data manipulation and mining activities. A strong collaboration between these two roles is essential, as the integration of domain knowledge with data is crucial for the success of a PHM project. In this particular project, however, the roles are more intertwined. My background as an aerospace engineer specialized in aircraft systems positions me as the SME, while I also engage in data manipulation and processing. Consequently, I find myself serving both as the SME and the data scientist in this project. This unique situation poses me a privileged perspective and provides me with hybrid competencies in both domains.

Finally, once one starts working seriously with data, the phrase “Garbage in – Garbage out” becomes more meaningful every day. In the context of this work, I would rephrase it as: “A PHM framework is only as good as the quality and relevance of the data it is built on.” Often, equipment-level PHM systems are developed on equipment-level data. This consideration, which may seem trivial, actually hides a central pillar: diagnosis and prognosis can only be performed if the health status of the system is, at least in part, observable. This requirement is often intrinsically satisfied, since the most direct way to monitor equipment health is to measure its signals. However, this is not the case addressed in this thesis, which focuses instead on deriving condition insights for aerospace components using data available at the aircraft level.

These considerations connect directly to the question of how prepared the industrial sector really is to embrace PHM. On the one hand, vertical organisation of data and responsibilities – with information siloed by system, department, or supplier – often prevents the integration needed to reconstruct a coherent picture of asset health. On the other hand, horizontal organisation across the asset lifecycle – from design to operation, maintenance, and decommissioning, and across different fleets and operators – is still far from fully realised. PHM thrives when these vertical and horizontal dimensions are aligned: when data can flow securely and meaningfully between actors, when models can be updated with feedback from the field, and when operational decisions can be informed by a shared, consistent view of risk and degradation.

This thesis was written with that future in mind. It reflects both the enormous potential of PHM and its very concrete limitations: incomplete observability, fragmented data, legacy systems, and organisational inertia. At the same time, it shows that, even under these constraints, it is possible to extract value from operational data, design robust frameworks, and move a step closer to condition-based and risk-informed decisions. Looking ahead, I am convinced that PHM will increasingly become a cornerstone of industrial asset management, not as a stand-alone "add-on", but as an integral part of how systems are designed, certified, operated, and maintained. The hope behind this work is that the methods, lessons, and reflections presented here can contribute, even in a small way, to that transition, helping future engineers and practitioners navigate both the possibilities and the limitations of PHM with clarity, realism, and ambition.

Contents

List of Figures

List of Acronyms

AHM Aircraft Health Monitoring.

AJT Advanced Jet Trainer.

ALIS Autonomic Logistics Information System.

AOG Aircraft on Ground.

ARP Aerospace Recommended Practice.

BITE Built-In Test Equipment.

CBM Condition-Based Maintenance.

CF Cumulative Feature.

CI Condition Index.

CMT Clustered Mission Type.

DDV Direct Drive Valve.

DL Deep Learning.

EASA European Union Aviation Safety Agency.

EDA Exploratory Data Analysis.

EHA Electro-Hydraulic Actuator.

EMA Electro-Mechanical Actuator.

EMAR European Military Airworthiness Requirements.

ENAC Ente Nazionale per l'Aviazione Civile.

ER Equipment Register.

ETL Extract, Transform, Load.

FBW Fly-By-Wire.

FCC Flight Control Computer.

FCS Flight Control System.

FD Flight Duration.

FDI Failure (or Fault) Detection and Identification.

FDR Flight Data Recorder.

FH Flight Hour.

FP Flight Parameters.

GM Guidance Material.

HT Horizontal Tail.

HUMS Health and Usage Monitoring System.

IIoT Industrial Internet of Things.

IPR Intellectual Property Rights.

IVHM Integrated Vehicle Health Management.

LCA Life Cycle Assessment.

LCC Life Cycle Cost.

MAWA Military Airworthiness Authorities Forum.

MEA More Electric Aircraft.

MMEL Master Minimum Equipment List.

MRO Maintenance, Repair and Overhaul.

MSG Maintenance Steering Group.

MX Maintenance.

O&M Operations and Maintenance.

O&S Operations and Support.

OEM Original Equipment Manufacturer.

OH Operating Hour.

PAS Primary Actuation System.

PDF Probability Density Function.

PF Particle Filter.

PHM Prognostics and Health Management.

PIT Pilot Input Transducer.

PN Part Number.

PNRR Piano Nazionale di Ripresa e Resilienza.

PoF Probability of Failure.

RCM Reliability-Centered Maintenance.

RDT&E Research, Development, Test and Evaluation.

RUL Remaining Useful Life.

SAE Society of Automotive Engineers.

SHM Structural Health Monitoring.

SLR Systematic Literature Review.

SM Statistical Moment.

SN Serial Number.

SOM Self-Organizing Map.

STNR Signal-to-Noise Ratio.

TN Tail Number.

UAV Unmanned Aerial Vehicle.

UR Unscheduled Removal.

VHMS Vehicle Health Management System.

WP Work Package.

Chapter 1

Introduction

1.1 Introductory Statements and Research Aim

Over the past decade, industrial maintenance has undergone a structural transformation. The combination of enhanced sensing and increased computational capabilities combined with widespread connectivity is progressively shifting maintenance from reactive and schedule-driven practices towards data-enabled, condition-based and prognostics-oriented strategies.

Within this context, Prognostics and Health Management (PHM) has emerged as a key paradigm for exploiting operational data to assess asset health, predict future degradation, and support maintenance decision-making. However, the practical deployment of PHM in aviation is still uneven. Although modern systems and engines are often designed with health monitoring in mind, many operational fleets consist of legacy platforms and systems whose design did not anticipate data-driven PHM. Flight control actuators, and in particular electro-hydraulic actuators (EHAs), represent a prominent example: they are mission-critical components, subject to complex loads and failure modes, yet their in-service monitoring is frequently limited to basic Built-In Test Equipment (BITE) and fault logging.

The work presented in this thesis is situated precisely at this interface between PHM theory and industrial deployment for legacy aerospace systems. It focuses on an operational fleet of Advanced Jet Trainers (AJTs) and on a specific component of their Flight Control System (FCS): the electro-hydraulic actuator of the horizontal tailplane. The fleet operates under demanding usage profiles and stringent safety and availability requirements, while maintenance is governed by existing standards, regulatory constraints, and cost considerations. At the same time, a wealth of heterogeneous data – including Health and Usage Monitoring System (HUMS) flight data, flight hours, equipment registers, and unscheduled removal logs – is already available, albeit not yet systematically exploited for prognostics.

The overarching aim of the research is to design a PHM framework that can leverage these existing data sources to provide probabilistic health information and actionable maintenance support for the selected actuator, without requiring intrusive hardware modifications or disruptive changes to current maintenance practices. The approach is deliberately conservative: rather than pursuing complex black-box models detached from operational constraints, the work seeks to build a coherent chain that starts from the data actually collected in the fleet, extracts physically interpretable indicators of usage and degradation, and translates them into predictions and decision-oriented metrics that can be integrated into the operator's processes.

More specifically, the thesis is articulated into the following research objectives, organized into Work Packages (WP) (Figure 1.1):

- **WP1 – Context and requirements.** To characterize the industrial and operational context in which PHM for flight control actuators operate, including Industry 4.0 trends and state-of-the-art solutions. This work package has been accomplished with the journal paper [3] and the conference proceedings paper [4].

- **WP2 – PHM functional architecture.** To define a functional PHM architecture tailored to the selected case study, structured into clear layers for data management, data processing, and health management, and compatible with legacy fleets and existing HUMS infrastructure. This work package has been accomplished with the journal paper [1].

- **WP3 – Data management and ETL.** To develop a scalable data management and ETL workflow capable of ingesting, cleaning, and harmonizing heterogeneous sources (HUMS LOAD files, equipment registers, maintenance documentation) into a consistent data model suitable for analysis. This work package has been accomplished with the conference proceedings [5].

- **WP4 – Degradation indicators.** To engineer and validate degradation-sensitive features and flight-level indicators that capture the relationship be-

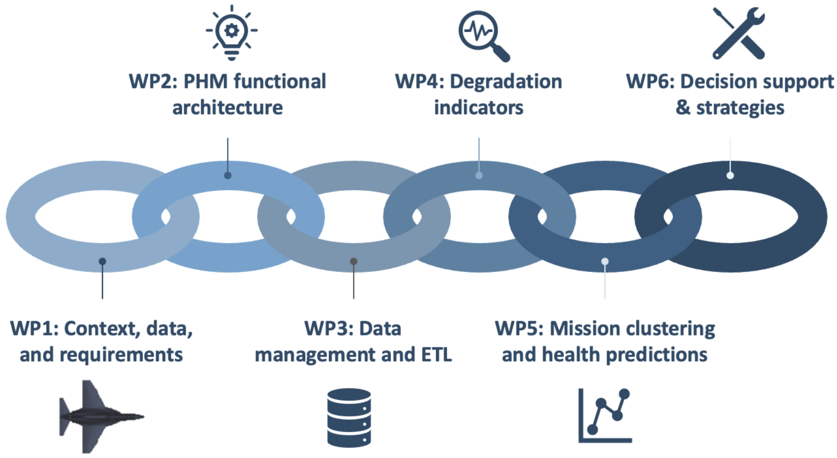

Fig. 1.1 Research objectives and work packages.

tween operational usage and actuator stress while remaining interpretable for domain experts. This work package has been accomplished with the conference proceedings [5, 6].

- **WP5 – Mission Clustering and Probabilistic health prediction.** To construct probabilistic health prediction models that exploit these features to estimate remaining life or risk proxies for the actuator, explicitly accounting for future mission profiles and fleet heterogeneity. This work package has been accomplished with the conference proceedings paper [7], which received a best paper award.
- **WP6 – Decision support and maintenance strategies.** To propose decision-oriented metrics and maintenance strategies, in both the short–mid term and the long term, that translate prognostic outputs into recommendations compatible with the operator's cost structure, operational constraints and risk posture. This work package has been briefly approached in the journal paper [1] and better expressed in the body of this thesis and future publications.

In line with this aim, the thesis is organized as follows. Chapter 1 sets the stage by introducing the broader industrial context (Industry 4.0 and aviation), clarifying what PHM is meant by and how it is implemented in practice, and discussing key change drivers such as costs, maintenance taxonomies, the regulatory framework, and existing standards. Then it reviews the relevant literature, provides a concise overview of flight controls and EHAs, and describes the selected AJT case study and its Flight Control System.

Chapter 2 develops the rationale for the project and presents the overall workflow adopted in the thesis, explaining the reasons for the selected PHM strategy and how the different modules interact. Chapter 3 formalizes the PHM functional architecture, structured into three layers: (i) Data Management, ETL and Pre-Processing, (ii) Data Processing, Features and Prognosis, and (iii) Health Management, Strategies and Support.

Chapters 4 to 6 then implement these layers in detail. Chapter 4 describes the Layer 1 modules devoted to data acquisition and pre-processing, including the handling of HUMS LOAD files, equipment registers, and technical documentation. Chapter 5 addresses Layer 2, covering feature engineering, flight clustering, probabilistic health predictions, and the computation of probability-of-failure and risk metrics, together with a critical discussion of performance and limitations. Chapter 6 focuses on Layer 3, where prognostic results are embedded in short- to mid-term and long-term maintenance methodologies, illustrating how the proposed framework can support health management decisions. Finally, Chapter 7 summarizes the main findings, discusses their implications, and outlines directions for future work.

This research has been carried out within an industrial PhD programme in collaboration with Leonardo S.p.A. and is co-funded by the company in the framework of the Italian National Recovery and Resilience Plan (PNRR), which implements the European Union's NextGenerationEU initiative to foster digitization, innovation and research. This context has strongly influenced the orientation of the work towards solutions that are not only methodologically rigorous but also realistically deployable and aligned with the strategic objectives of both the industrial partner and the wider European effort to modernize and digitalize critical infrastructures.

1.2 Industry 4.0 and Aviation

Maintenance can be defined as:

> *"The process of ensuring that a system continually performs its intended functions at its original level of reliability and safety"* [8]

"Any one or combination of overhaul, repair, inspection, replacement, modification or defect rectification of an aircraft or component" [9]

As the global aircraft industry continues to evolve, the maintenance sector stands right at the crossroad of opportunity and challenge.
The opportunity of providing customers with cutting edge and integrated digital threads and the challenge of data set quality and integrity.
The opportunity of forestalling equipment health status and down times and the challenge of dealing with siloed and unsystematic data management.
The opportunity of enhanced fleet availability and readiness and the challenge of conceiving cost-effective and dependable actionable data-driven strategies.

In an era where data drives decisions, the convergence of technology and data analytics is reshaping how we approach asset management and system monitoring. Our society is becoming increasingly more data-centric and every new system or equipment around us is being designed to log and save operational data which can be used to improve usability and operation. However, it must not be forgotten that does exist a wide range of already operational in-service systems which were designed when the idea of using data to enhance the platform operability was still not consolidated. All these systems and platforms can be brought on-line to leverage the benefits of data driven operative optimization if forward looking projects are approached.

The concepts of Industry 4.0 and the paradigm of Industrial Internet of Things (IIoT) are driving the industrial sector towards an ever increasing system digitization and sensorization. As a result, more and more data are saved and can provide useful insights into machine operations and potential future malfunctions.
A possible outcome of data analytics integration into management decision-making stands in the development of intelligent and adaptive maintenance strategies which can support the equipment only when needed, thus optimizing availability and safety [10].

A wide range of research papers, technical reports, and business insights [11] is showing that the rapidly evolving aircraft Maintenance, Repair, and Overhaul (MRO) sector is a field that will be subject to revolutionizing technologies in the next decades. The advent of Industry 4.0 has shown the enormous value of data as a tool to better understand the behavior of equipment and has placed data itself as the center of a new era. In fact, if data have been obviously already incorporated in the design, engineering and manufacturing stages of an aircraft life cycle, not the same can be said for the operating and service phase. A 2020 report from the consulting firm ICF [12] shows that while the integration of a digital thread in the first phases of design and manufacturing has reached a maturity level of more than 85 per cent, the integration of data driven solutions for the operation and service phases is as low as 30-40 per cent, revealing a consistent gap in how data can be exploited.

This direction is also reflected in Leonardo's Industrial Plan update (2025–2029), which explicitly positions digitalization as a core enabler of competitiveness [13]. In parallel, Leonardo frames the transformation toward multi-domain operations as the evolution from domain-specific solutions to an integrated approach, where digitalization and cyber protection by-design constitute the connecting thread across products, systems, and processes [14]. Together, these strategic elements motivate the adoption of data-centric architectures and service-oriented solutions, coherently aligning with the Industry 4.0/IIoT paradigm and the development of digitally enabled, adaptive MRO practices.

The direction of bringing asset online and the creation of a digital continuum which could enable a data centered MRO is currently a hot topic in the aviation service providers community. According to a 2019 Cap Gemini survey conducted among aircraft operators [12] on "Digital Aviation in MRO", 87.4 per cent of the respondents expressed that their existing MRO IT solutions need to be improved. In contrast, the report also indicated that 76.9 per cent of the aircraft operators surveyed recognized the potential value that could be derived if Original Equipment Manufacturers (OEMs) were to address and bridge digital thread gaps. In addition, Al-Kaabi et al. [15] states that essential MRO activities, such as line Maintenance, are typically managed in-house, whereas activities with lower demand at the airline level, such as equipment maintenance, are often outsourced, underscoring the importance of a solid ground base on which to build a data framework. In essence, OEMs are viewed as essential entities in providing a digital thread solution and closing the existing digital thread gap.

1.3 PHM: What, Where, and How

Prognostics and Health Management (PHM) is a discipline that involves the monitoring of asset data and condition to predict and prevent potential failures. Taking inspiration from space oriented technologies developed in the mid-1990s, PHM technology has been first introduced in the early 2000s along with the development of the JSF Autonomous Logistics Integrated System (ALIS) [16, 17]. In this regard, PHM strategies were promptly identified as the primary enabling technologies for the development of enlightened maintenance frameworks. Since then, the PHM discipline has been applied to a wide range of engineering sectors, once again proving the solid and robust foundation and idea of the JSF project. In more than 20 years, PHM strategies flourished and their applications sparked interest not only within the development of new products but also in view of applying PHM principles to already operational legacy assets.

PHM is often described using the notion of functional layers. Many studies identify three main layers that, once combined, form the core of PHM methodologies: the diagnosis layer, the prognosis layer, and the health management layer [18]. The diagnostic layer is responsible for the tasks of detecting, isolating, and quantifying faults or failures within the system. This layer plays an important role in the early identification of issues, allowing timely interventions that can prevent further complications. The overall process is often referred to Failure (or Fault) Detection and Identification (FDI). In this sense, it is important to clarify the meaning of the terms fault and failure. In the context of this thesis, as often accepted in the PHM community, a fault is intended as a malfunction, which, however, does not preclude the overall functioning of the system. However, a fault can become a failure if it is not addressed in time in the so-called fault-to-failure process. As a result, a failure does affects the operation and safety of the system.

Once a fault is identified, the prognostic layer is dedicated to assessing the performance and integrity of the system and predicting its Remaining Useful Life (RUL). By analyzing various parameters and historical data, this layer provides insight into how long the system can continue to operate effectively before maintenance or replacement is necessary. Finally, the health management layer utilizes the information gathered from both the diagnostic and prognostic layers to optimize overall asset operations and mission readiness in various ways, ultimately streamlining processes and reducing costs: maintenance optimization, spare parts management, etc.

The creation of enlightened and forward-thinking maintenance strategies goes hand in hand with the development of PHM strategies. PHM can transform traditional maintenance practices by implementing a proactive framework that fosters more reliable and cost-effective management of critical systems. In this sense, PHM is an essential foundation upon which innovative maintenance strategies can be based.

1.3.1 HUMS as a PHM Enabler

Aircraft, as many other complex systems, are nowadays equipped with logging and monitoring systems which provide enhanced diagnostics and monitoring capabilities. The name of the system may vary according to the specific platform or proprietary names; however, one of the first systems developed for these purposes is called Health Usage Monitoring System (HUMS). A non exhaustive list of the possible systems is here reported. There are several alternatives, each offering unique capabilities for monitoring and analyzing aircraft health: Aircraft Health Monitoring (AHM) [19], Aircraft Condition Monitoring Systems (ACMS) [20], Vehicle Health Monitoring Systems (VHMS) [21], Integrated Vehicle Health Management (IVHM) [22], etc. The name, design and specific applications of HUMS-like systems may vary drastically between applications, but the underlying principle remains unchanged: collecting and analyzing data from various aircraft components and systems to assess their health condition. These systems can include software-based solutions that monitor sensors all around the aircraft, capturing thousands of parameters from multiple sensors, allowing operators to customize their data acquisition. These systems prioritize the ability to monitor conditions and identify possible anomalies, faults, or failures. In the 1990s, HUMSs were developed initially for rotary wing applications where vibration monitoring is of central importance for the health of the gearbox, mechanical transmission, and components in general [23]. As years passed, HUMSs began to be employed in several contexts as they demonstrated their potential to detect faults, improve safety, and increase availability by reducing unexpected downtimes [24]. HUMS (or analogous) systems can be viewed as a ground base to design PHM applications in legacy aircraft, since they provide the data acquisition and logging module [25]. For instance, data from a HUMS are used in the development of maintainability improvement processes for Electro-Mechanical Actuators (EMAs) in [26].

More often than not, when dealing with in-service platforms, PHM system designers are faced with the challenge of working with a system that, despite being envisioned to perform similar tasks, is not conceived for PHM. In other words, data sampling, sensors, and logging rate are often optimized for basic diagnostic tasks in an already outdated and old-fashioned system [25]. Nevertheless, HUMS are often the only source of time-series data available for the development of PHM frameworks. Kappas and Frith in [25] addressed the exact challenge of trying to connect the dots between HUMS and PHM and highlighting the evolution of HUMS from basic data acquisition systems, primarily hardware-focused, to more sophisticated, software-driven integrated health management systems.

1.4 Change Drivers: Costs

The effort of aerospace industry stakeholders towards the development of integrated data solutions is supported by the magnitude and direction of overall economic trends, since the market size for the aerospace MRO sector is estimated at USD 127.98 billion USD in 2024 [27]. The sole military MRO sector is valued at around 37 billion USD in 2024 with a Coumpound Aggregated Growth Rate (CAGR) of 2.84 per cent in the years 2024-2033 [28]. The Military Aviation Market size alone is estimated at 55.63 billion USD in 2024, and is expected to reach 75.52 billion USD by 2030, growing at a CAGR of 5.23 per cent during the forecast period (2024-2030) [29]. The market outlook for the next ten years is positive, also because the military MRO sector has not been deeply influenced by the COVID-19 crisis, unlike its civil counterpart and recent news highlighted the plan to invest more in defense sectors. Demand is influenced by an expanding and/or aging military fleet and, on the other hand, by shrinking defense budgets [30]. These drivers guarantee a steady growth as defense administrations are keen to maximize the value and lifespan of their assets. With a positive market outlook and a wide range of business opportunities, the desire to streamline processes and reduce costs is of paramount importance. In fact, MRO optimizations could drive the change by reducing a big slice of aircraft related costs.

From a life-cycle cost perspective, aircraft-related costs are commonly organized into broad categories that reflect the main phases of system acquisition, operation, and end-of-life management. In earlier literature, the discussion was often centered on the three major categories of (i) Research, Development, Test and Evaluation (RDT&E), (ii) Procurement, and (iii) Operations and Support (O&S), with the latter also often referred to as Operations and Maintenance (O&M) [31]. In more recent cost-estimating frameworks, disposal is explicitly identified as an additional life-cycle cost category, reflecting a more complete Life Cycle Assessment (LCA) perspective [32]. In practice, however, disposal costs are typically far lower than the other categories, and the first three remain the dominant contributors to the overall life-cycle cost [32].

For many years, cost analysts have been focusing mainly on the first two categories and did not consider the impact of O&S costs. However, the practical importance of accurate O&S estimates is manifesting itself in programmatic decisions and it is becoming day to day more important for decision makers [33] and for the allocation of long-term maintenance budgets. Especially in the aerospace sector, operating costs have become increasingly dominant within the overall life-cycle cost structure, highlighting the need for innovative and forward-looking projects that specifically target these expenditures in order to improve the economic sustainability of aircraft operations. Historically, this observation has often been summarized by the so-called "70:30 rule", which describes the approximate ratio between Operations and Support (O&S) costs and the combined costs of Research, Development, Test and Evaluation (RDT&E) and procurement. This rule was already highlighted in the late 1990s and provided a useful benchmark for cost estimations in the defense aerospace sector [31]. The concept has since been widely referenced in policy documents and strategic analyses (e.g., reports from the U.S. Government Accountability Office), emphasizing how the majority of a platform's life-cycle cost is incurred after the system enters service, rather than during its development and acquisition phases [34]. Subsequent studies have refined this rule, suggesting that the proportion of O&S costs may be slightly lower—often closer to 60% of the total Life Cycle Cost (LCC) rather than 70%—depending on the platform type and operational context [35]. Nevertheless, the underlying principle remains robust: O&S costs represent the dominant share of life-cycle expenditures for most complex engineered systems, particularly in the aerospace domain. Recent analyses confirm these trends. The 2025 Operating and Support Cost-Estimating Guide published by the U.S. Department of Defense provides updated empirical evidence across several classes of defense systems [32]. Fixed-wing aircraft, rotary-wing aircraft, and UAV platforms exhibit O&S shares ranging approximately between 58% and 66% of total life-cycle cost, reinforcing the central role of operational costs in aerospace programs.

An analysis of O&S costs reveals several noteworthy insights. A 2015 RAND Corporation report highlights that, among the US Air Force and Navy Aircraft O&S Cost in the FY 2010–2013, pure maintenance costs accounted for the 33 and 40 per cent for USAF and USN respectively. A 2018 Air Force Institute of Technology

(AFIT) study [34] showed that, for a fighter/attack category aircraft, the sum of Unit Manpower, Unit Operations and Maintenance reaches an astonishing 87.42 per cent of all O&S costs (33.78, 17.68 & 35.96 respectively). Pure maintenance costs are estimated around 35 per cent, in line with the aforementioned other reports. In the civil sector, the magnitude of MRO is underscored by its representation of approximately 12-15 percent of an airline's O&S costs [36]. In 2013, annual expenditures in this sector were estimated at US $ 50 billion [37], supporting the employment of around 480,000 individuals globally [38].

In conclusion, if, on the one hand, the military MRO market shows promising growth in the next year, fueled by the need to keep increasingly aging fleets available, on the other hand, the need for cost-effective practices and optimized procedures in aircraft operations drives the research towards advanced maintenance solutions. With Aircraft-On-Ground hourly costs reaching several thousands of US dollars [39], high platform availability is required. This once again underlines the economic benefits that the integration of progressive maintenance activities would bring [40, 41]. It is worth noting that a 2023 report from Market Research Future indicated that the generic predictive maintenance market, which includes various industrial sectors such as aerospace, was valued at USD 17.3 billion in 2021. The sole Predictive Maintenance market industry alone is expected to increase from USD 21.83 billion in 2022 to USD 111.30 billion by 2030, demonstrating a CAGR of 26.20% during the forecast period of 2024 to 2030. In this complex and multifaceted environment, a very promising approach is indeed represented by the adoption of enlightened maintenance strategies, which promise to potentially change stable maintenance paradigms and improve asset performance, availability, and safety by incorporating asset usage data into maintenance decision making [42].

1.5 Maintenance Taxonomy

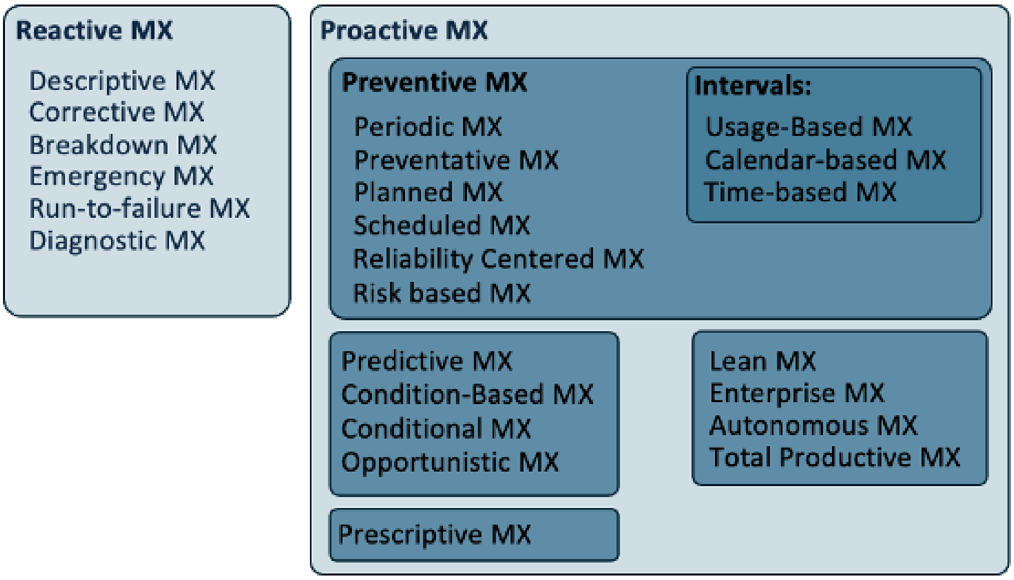

Fig. 1.2 Graphical representation of the different Maintenance (MX) strategies in a Venn diagram fashion

The realm of maintenance approaches is vast as an extensive range of definitions has been created during the years: Reactive, Breakdown, Corrective, Run-to-failure, Emergency, Descriptive, Diagnostic, Preventive, Preventative, Scheduled, Planned, Time-based, Calendar-based, Periodic, Risk based, Reliability Centered, Usage-Based, Predictive, Conditional, Condition-Based, Opportunistic, Prescriptive, Enterprise, Lean, Total Productive or even Autonomous maintenance. A graphical representation of the relationships and hierarchy between these maintenance approaches is reported in Figure 1.2 and adressed in the next paragraph. However, maintenance strategies can be mainly divided into two broad categories: Reactive maintenance and Proactive maintenance.

1. When Reactive maintenance is adopted, assets and equipment are allowed to run until failure before maintenance is performed. Maintenance is conducted only when equipment failure occurs and actions are triggered by failed equipment tests, system malfunctions, faults/failures, etc. It is the "if it isn't broken, don't fix it" approach. If this maintenance approach is selected, the focus of maintenance actions is on restoring the equipment to its normal operating state

as quickly as possible, rather than preventing it. The run-to-failure approach is adopted for non safety-critical, disposable, durable, or non-maintainable, low capitalization or low-cost assets that do not significantly impact aircraft operations (e.g., as indicated in the aircraft Master Minimum Equipment List (MMEL)). Reactive maintenance includes Breakdown maintenance, Corrective maintenance, Run-to-failure maintenance and it occurs when the equipment breaks down requiring a corrective action. It is the most basic kind of maintenance and it may require an unscheduled maintenance action depending on the importance of the broken equipment. In this way, reactive maintenance includes the Descriptive or Diagnostic maintenance category as well, since a fault or failure diagnosis or description is required to make the equipment operational again. Emergency maintenance is another synonym of Reactive maintenance which is carried out in emergency conditions and involves safety critical or flight critical equipment.

2. Proactive maintenance involves performing maintenance tasks before any major breakdowns or failures occur. Various criteria can trigger the need for proactive maintenance activities, resulting in several types of proactive maintenance strategies based on the different criteria. The term Proactive maintenance often refers to a wide plethora of solutions which has evolved over time to implement innovative maintenance approaches.

    - Scheduled maintenance, Planned maintenance, Time-based maintenance, Calendar-based maintenance, Periodic maintenance can be grouped in the Preventive or Preventative maintenance category. In particular, the intervals can be based on:

        - time intervals
        - calendar intervals
        - cycles (Usage Based)

      The maintenance planning can also be risk based maintenance or the so called Reliability Centered maintenance (RCM). Despite being innovative, these approaches are static and the maintenance plan does not change during the aircraft life.

    - On the other hand, Predictive maintenance, Conditional maintenance, Condition-Based maintenance are custom based maintenance strategies which follow the aircraft life adapting to the condition of the asset being monitored with some differences. Conditional or Condition-based maintenance is an approach based on the monitoring of the actual condition of an asset to proactively decide what maintenance needs to be done according to the actual state of the system. The objective of Condition-based maintenance is to conduct maintenance activities based on the current condition of the equipment, rather than adhering to a predetermined schedule or addressing failures. The difference between Predictive maintenance and Condition-based maintenance is shady and it can be interpreted in various ways. Some scholars highlight how Predictive maintenance is a step forward with respect to Condition-based maintenance since it tries to predict the future state of the system while Condition-Based maintenance only detect abnormal behavior on monitored data and without prognosis, as such, often the difference between these two maintenance strategies is timing. However, the differences are very subtle and the terms Condition-based maintenance and Predictive maintenance are often used interchangeably in literature underscoring an innovative maintenance strategies that aim to reduce down times exploiting historical and operational monitoring data.

- Rising in predictive power, the next on the list is opportunistic maintenance. Opportunistic maintenance is a maintenance approach that involves the strategic replacement of equipment items or components during planned or unplanned system shutdowns. This approach leverages the availability of maintenance resources already on site to optimize efficiency. In other words, each asset stoppage can be seen as an opportunity for preventive maintenance on other equipment and, by grouping up maintenance tasks when the asset is already down, maintenance can be effectively optimized.
- Then, prescriptive maintenance is a proactive maintenance strategy that exploits asset monitoring data to determine and recommend required maintenance actions on equipment, anticipate asset maintenance requirements thus providing information about how to delay or entirely delete equipment impromptu failures.

- Finally, if maintenance is seen as integrated into a higher level structure, we can talk about enterprise maintenance, lean maintenance, total productive maintenance or even autonomous maintenance.

1.6 Regulatory framework

In the context of European civil aviation, airworthiness regulations are governed by common rules established by the European Aviation Safety Agency, pursuant to Regulation (EC) No 216/2008 of the European Parliament and of the Council. This framework includes two key Implementing Rules: Commission Regulation (EU) No 748/2012, which pertains to initial airworthiness, and Commission Regulation (EU) No 1321/2014, which addresses continuing airworthiness. These documents are mandatory and must be adhered to by all civil aviation organizations and individuals operating within Europe. European Civil Aviation Authorities (such as ENAC in Italy) are required to implement these rules into their national regulations without deviation.

For national sovereignty reasons, European Regulation (EC) No 216/2008 does not pertain to products, parts, appliances, personnel, or organizations engaged in military, customs, police, search and rescue, firefighting, coastguard, or similar activities or services. These activities and services are governed solely by national legislation and the responsibility rests entirely with the Member States. Nonetheless, Member States are encouraged to ensure that these activities and services consider the objectives of these regulations to the extent practicable. For instance, within EASA countries, a civil MRO must be approved in accordance with EASA Part 145 contained in Regulation 1321/2014. However, this does not apply to military MROs.

As a result, participating Member States have established their own national systems for managing the airworthiness of military aircraft. To promote standardization and coordination in this area, the European Military Airworthiness Requirements (EMARs) are being developed. Specifically, the Military Airworthiness Authorities (MAWA) Forum, with the support of the European Defense Agency (EDA), has been established to create harmonized military airworthiness processes among its participating Member States. This initiative will facilitate the development of a unified set of EMARs, along with accompanying Acceptable Means of Compliance (AMC) and Guidance Material (GM) that can be integrated into the national military airworthiness regulations of all participating Member States, drawing upon best practices from civil aviation regulations. The benefit of a stronger coherence between Member States is of utmost importance for a variety of reasons: first of all, a better platform interchangeability capabilities along with potential savings in terms of time, cost, effort and better Pooling and Sharing opportunities in the development of new platforms. In particular, the EMARs pertaining to Continuing Airworthiness are as follows:

- EMAR M (comparable to EASA PART M): Outlines the necessary measures to ensure the continued airworthiness of aircraft, along with the conditions that must be fulfilled by individuals or organizations responsible for managing ongoing airworthiness.
- EMAR 145 (comparable to EASA PART 145): Defines the criteria that organizations must satisfy to obtain or maintain approval for the maintenance of aircraft and their components.
- EMAR 66 (comparable to EASA PART 66): Specifies the educational and training requirements for personnel involved in aircraft maintenance.
- EMAR 147 (comparable to EASA PART 147): Details the standards that organizations must meet to gain approval to conduct aircraft maintenance training and examinations.

1.7 Existing Standards and Procedures

Literature studies show that, by design, a good PHM system should be fit for purpose, Integrated, Robust, Reliable, Relevant, User Friendly, scaleable, and Cyber secure [25]. Other sources, such as the work presented by Jardine et al. [18], provides a higher perspective of a Condition-Based Maintenance (CBM) system, stating that CBM requires three steps: Data Acquisition, Data Processing, and Maintenance Decision Making as already mentioned in Section 1.3. These concise and abstract phases can be further expanded in operational guidelines. To design a scalable, well-thought and effective PHM system, a comprehensive review of key PHM standards in the aerospace sector, as well as in other relevant fields, has been accurately studied

[43]. As a result, numerous standards and guidelines available for the development of PHM systems have been identified:

- ISO 17359:2018 "Condition monitoring and diagnostics of machines — General guidelines" [44]. This document provides guidelines for establishing a condition monitoring program for machines and includes references to related standards pertinent to this process.
- OSA-EAI [45]. The Open System Architecture for Enterprise Application Integration (OSA-EAI) specification provides (i) an information exchange standard to facilitate sharing asset registry, condition, maintenance, and reliability data between enterprise systems, and (ii) a relational database model for the storage of this asset information.
- ISO 13374:2003 "Condition monitoring and diagnostics of machines — Data processing, communication, and presentation" [46]. This resource is more focused on the information technology part and may provide useful indication on the data and information management in a PHM enlightened perspective.
- ADS-79E-HDBK [47]. This document offers the privileged perspective of the United States Army towards the development of CBM systems for US Army aircraft.
- OSA-CBM. OSA-CBM was developed by, principally, an industry consortium under the DUST (Dual Use Science and Technology) program and further managed by the MIMOSA group [48], who acts as a repository for the standard. The MIMOSA group is a non profit organization aimed at enhancing and encouraging the adoption of open standards enabling physical asset life-cycle management. This is more a recommended architecture specification rather than a strict standard or procedure; nevertheless, it offers useful information for transferring data and knowledge in a CBM system.
- SAE ARP 6883:2019 [49]. This Aerospace Recommended Practice (ARP) developed by the Society of Automotive Engineers (SAE) provides information on the development of IVHM strategies for complex systems.

In addition, it should be noted that MSG-3 (Maintenance Steering Group 3) is the industry standard methodology developed over the years to define scheduled maintenance tasks and intervals for aircraft systems. Traditionally focused on reliability and safety, MSG-3 provides a structured, logic-based framework to identify preventive maintenance actions. In recent years, however, new developments have aimed to extend and modernize MSG-3 to incorporate PHM concepts, enabling data-driven and condition-based maintenance within the same decision logic. The SAE IVHM Committee HM-1 is actively working towards this integration, developing guidelines and standards to align IVHM methodologies with existing MSG-3 logic and to support a gradual evolution toward predictive, evidence-based maintenance programs.

These documents, along with the research group experience in managing similar projects, have been crucial to conceiving an effective workflow, which will be presented in Chapter 2.3.

1.8 Literature Review and Related Works

As stated in Section 1.3 the aerospace sector has been the first field where PHM was applied and, since then, many aircraft subsystems have been the focus of PHM research. However, hydraulically powered flight control systems have not yet received the same interest as other subsystems, especially by using operational data. In fact, while engines, batteries, landing gear brakes, and Environmental Control Systems (ECS) have been addressed in several PHM studies [50–57], the FCS still lacks a solid PHM base. Some studies focus on PHM for Electro Mechanical Actuators (EMAs) following the paradigm shift of More Electric Aircraft (MEA) [58, 59], but the literature panorama for EHA is scarce and scattered. This fact may be explained from different points of view:

1. EMAs are projected to replace EHAs in the coming years. Consequently, many research centers are directing their efforts toward PHM, anticipating it to be the primary solution in the future. However, it is important to recognize that EHAs are currently used in all commercial and military aircraft.

2. EMAs are extensively sensorized, while EHAs are often lacking in sensors. The electric operation of EMAs facilitates the easier monitoring and recording of performance signals. In contrast, EHAs pose challenges for sensor integration, as OEMs often prefer to avoid adding sensors to preserve actuator

logistic reliability. Consequently, it is generally more straightforward to track health indicators from EMAs, such as phase currents, compared to monitoring signals from EHAs.

3. PHM was established at a time when EHAs were already widely adopted, while EMAs are currently emerging as a technology for primary flight control systems.

4. PHM is often proposed as a solution to support the adoption of EMAs in safety critical applications (i.e., primary flight controls) because their overall reliability is, at the current stage, not enough to permit their use for primary flight controls. In this sense, PHM could be employed to back EMAs up for their acceptance in flight critical domains. This is not the case for EHA, which are extremely safe, rugged, time-tested and do not suffer from single points of failure (unlike EMAs) [60, 61]. PHM for EHAs would not be implemented to potentially support a certification phase but only to increase the platform availability, improve readiness, reduce costs, and enhance the performance.

5. There is a consistent lack of publicly available EHAs data sets in operations due to IP restrictions and proprietary policies.

These reasons explain the lack of a substantial body of literature involving PHM strategies for EHAs, especially with operational data, even if PHM strategies in general are mature and well known. An extensive Systematic Literature Review (SLR) has been carried out to identify potential similar studies. The results of the SLR have been explained in detail in a journal paper [3]. As already mentioned, the results are scattered and scarce; however, they can be categorized based on three main groups.

- Category 1: This category includes studies that utilize operational and fleet data to evaluate the health status of various systems, although these systems are not directly related to the specific case study (i.e., they do not study EHAs). These studies often provide valuable information on general health monitoring techniques that can be applied to different types of equipment that have operational data available.

- Category 2: This category includes research that specifically targets EHAs but relies on equipment-level data that are typically difficult to obtain in real-world operational settings. Such studies may involve the use of additional sensors, controlled laboratory tests, or advanced modeling techniques to gather comprehensive data on EHA performance. While these insights are important for modeling specific EHA fault modes, behaviors, or degradations, the methodologies in this group are hardly applicable to real-life data.

- Category 3: Publications in this category focus on the development of PHM methodologies specifically for EHAs, utilizing both operational and fleet data. These studies aim to create practical frameworks for monitoring and predicting the health of EHAs in real-time scenarios.

Each group will be further analyzed in the next paragraphs.

Three studies focused on PHM using operational data that could be compared with available data from the project. In the first category, two papers published in connection with the 2014 PHM Society Data Challenge have explored a similar issue, highlighting the challenges associated with sparse and fragmented data. Although the equipment studied was not directly related to EHAs, the available data in an industrial context shared certain similarities with that of the research approached in this thesis. Rezvanizaniani et al. [62] emphasize the challenges of managing logistic and operational data from various sources, proposing solutions grounded in probabilistic risk assessments. Their strategy relies on statistical methods, utilizing the extensive number of Unscheduled Removals (URs) and detailed part removal data present in the dataset. In contrast, Kim et al. [63], while working with the same dataset, adopted a completely different statistical approach based on an ensemble technique that combines parts lifespan calculations with various usage classification methods. Unfortunately, these methodologies developed are not directly applicable to our research project because of the significantly lower volume of available data and the lack of precise indications regarding failure modes. When examining the aerospace sector, a Condition-Based Maintenance (CBM) framework applied to the propeller and wheel brake assembly of a C-130J aircraft is discussed by Schoenmakers in [64], using fleet and operational data. However, the results of applying several methodologies were limited. This once again underscores a critical point: Collecting more comprehensive data is essential to make reliable and informed decisions. The need for robust data collection methods is paramount in improving predictive maintenance strategies and ensuring the possibility of developing PHM strategies.

Another perspective is given by the second group of studies, those that specifically focus on EHAs but use data that is often challenging or impossible to obtain in real-world operational settings. Some of these approaches focus on the health of the overall EHAs as a whole, while other studies provide results for specific failure modes, specific components inside the EHAs (e.g. seals), or specific degradation patterns. These studies offer a range of solutions albeit focused on individual components. A non-exhaustive list of the major contributions is here reported. For instance, Mi and Huang [65] and Byington et al. [66] have specifically addressed servo valves, while Shanbhag et al. [67] conducted research on cylinders. Additionally, Bertolino et al. [68] present methodologies for detecting leakages, and Chao et al. [69] focus their studies on piston pumps. Bacci et al. [70] present a multi-physics model of a faulty rod-end for EHAs, whereas Vianna and Malere [71] outline a methodology for detecting hydraulic system leaks and providing service recommendations to mitigate Aircraft On Ground (AOG) risks. Their strategy utilizes three onboard sensors-a pressure transducer, a temperature transducer, and a quantity gage—whose data are recorded in the Flight Data Recorder (FDR) for a fleet of EMBRAER regional jets. Zong et al. [72] have created a real-time monitoring system for the actuator mechanism of an aileron. This study is framed within Structural Health Monitoring (SHM) strategies and emphasizes the comparison of dynamic responses between nominal and non-nominal systems through the use of finite element software and simulations. The studies reported above illustrate the significant body of research focused on actuator components or specific failure modes, highlighting the extensive literature available on this subject. These research works reveal a notable gap in addressing PHM at the Electro-Hydraulic Actuator (EHA) level, which, to the authors' best knowledge, is approached by only a limited number of research groups. Research by Liu et al. [73] discusses the predictions of performance degradation in EHAs using advanced ML techniques, including the Elman neural network observer, support vector regression (SVR), and the Gaussian Mixture Model (GMM). Additionally, studies conducted by Soudbakhsh and Annaswamy [74], as well as findings from Lu et al. [75], illustrate significant progress in fault detection methodologies and health monitoring strategies. Kordestani et al. [76] introduced a modular hybrid fault prognosis method that combines distributed neural networks with a recursive Bayesian algorithm. Furthermore, Cui et al. [77] proposed a hybrid approach that employs the Nonlinear Wiener Process (NWP) algorithm for the physics-based components and a data-driven Echo-State Network (ESN) for the data-driven aspects.

Numerous methodologies integrate an advanced data processing technique known as Particle Filtering (PF), a Bayesian approach employed for fault detection and Remaining Useful Life (RUL) estimation [78–80]. Particle filters (PFs) can be used as the basis to build probabilistic failure prognostic algorithms that estimate the future health state of a system based on degradation models and measurements [78–80]. Specifically, prognostic algorithms based on PF monitor and predict the system's state of health by analyzing the probability density function of multiple weighted samples, known as "particles," that evolve over time. The widespread success of PFs can be attributed to their ability to incorporate uncertainty into the estimation process, making them an optimal choice for providing estimates with associated levels of uncertainty and for their effectiveness in handling nonlinear and non-Gaussian distributed data. In the context of EHAs, PF has been utilized by Mornacchi et al. [81] and Autin et al. [82]. Variations of PF have also been explored by Guo and Sui [83], who implemented the Minimum Hellinger Distance technique alongside PF applications for EHAs. Finally, Macaluso and Jacazio have explored the extraction of relevant features during pre-flight checks and from in-flight data to enhance monitoring capabilities [84]. Although not comprehensive, this analysis offers a snapshot of the advancements and possibilities in PHM for specific failure modes and specific EHA components. Although these works offer valuable information, their applicability to existing legacy systems can be limited due to insufficient monitoring of low-level subsystem data and the lack of recorded control signals within the FCC loop. Once again, it is essential to note that these strategies require actuator-level data to effectively assess the health status of the equipment. In this context, it is important to recognize that a substantial number of studies concentrate on test bench approaches and synthetic datasets or focus on monitoring specific fault modes [84].

Lastly, only one record has been identified within the third category. Kannemans and Jentink [85] focus on the flight control actuator of a HT in a fighter aircraft, applying their methodology to data obtained from a Dutch Air Force F-16 Falcon FACE measurement system [86]. This particular research is performed by the Dutch PHM Consortium (DPC) and shares a lot of similarities with the thesis research project, as in both studies the aim is the estimation of the actuator usage leveraging operational data from a military jet. The main rationale is that the actuator rod displacement and actuator force are directly linked to control surface deflection and hinge moment. In this way, their research aimed to extract actuator usage information from hinge moment data via simulation and sensors already mounted on the aircraft as well as some test sensors added on purpose. However, the results were somewhat limited and there have been no further updates on the progress of this research project. In addition, the data used in the context of this article are not available for this research thesis. This highlights the challenges faced in utilizing operational data for in-depth analysis of flight control systems.

In conclusion, the detailed literature review indicates a relatively limited body of research in this area, underscoring the need for further investigation [3]. It should be noted that this research gap is particularly pronounced given the critical role that EHAs play in controlling primary flight control surfaces in all modern commercial aircraft and the majority military aircraft. The lack of comprehensive studies is especially striking considering the importance of these systems to aviation safety and performance and justifies further research.

1.9 Flight Controls and Electro-Hydraulic Actuators

Aircraft are often seen as interconnected systems of systems, whose interactions are essential for safe and reliable flight operations. The Flight Control System (FCS) is one of the most critical ones, and its aim is to transmit flight control commands from the cockpit or other sources, such as autopilots, to the appropriate actuators. This system is crucial in creating the forces and torques that set the position, speed, and attitude of the aircraft. To effectively move the aerodynamic surfaces and provide the necessary forces and torques, the FCS relies on actuators such as EHAs or EMAs.

Flight controls are often divided into primary and secondary flight controls.

- Primary flight controls are critical for the safe operation of an aircraft and, in a conventional configuration, include the ailerons, elevator (or stabilator), and rudder. These controls are flight critical systems and they are essential to control aircraft position, speed, and attitude.
- In contrast, secondary flight controls are employed to enhance performance and reduce the pilot workload. They often include devices such as flaps, slats, and trim systems. While losing secondary controls may complicate flying, it does not put the control of the aircraft in jeopardy except in very

specific asymmetry conditions. In some cases, slaps may also be employed continuously during flight to provide additional controllability. In this case, they are often called Leading Edge Flap Actuation Systems (LEFAS) and are used, among other high performance aircraft, in the Advanced Jet Trainer (AJT) under study, as better explained in Section 1.10.1.

- Finally, spoilers are flight control surfaces located on aircraft wing tip, serving two primary functions: enhancing roll control and acting as speed brakes. As a result, they can be classified as primary or secondary flight controls depending on the actual intended and designed use.

Today, the main choice for primary flight control actuators falls on EHAs due to their proved reliability, extreme ruggedness, and high power density. EHAs are therefore integrated within the aircraft hydraulic system. It has to be noted that, as already mentioned, in the last 20 years, a new paradigm is shifting the interest towards EMAs which, in the future, could be employed in primary flight control actuation systems. Currently, the substitution of EMAs in place of time-proven EHAs is still restricted by reliability issues and lack of enough data to allow certification for critical primary flight controls for safety [87]. As a result, to the best of the authors' knowledge, EMAs are employed in primary flight controls only in Unmanned Aerial Vehicles (UAV) or landing gear braking systems [88]. On the other hand, secondary flight controls are already often powered by EMAs, since safety requirements are less strict for this kind of flight controls.

This thesis is focusing on primary flight controls of a conventional aircraft and, as a consequence, the focus is on EHAs. EHAs are sophisticated pieces of equipment that are widely used not only in aerospace applications but also in various industrial sectors because of their robustness, reliability, and high power density. In aerospace contexts, where safety is dictated by many regulations, EHAs are typically configured in an active-standby tandem fashion. This configuration allows the standby cylinder to dampen external oscillations and vibrations, enhancing overall system stability and intervene in case of a failure on the other actuator chamber.

Describing a generic architecture for EHAs can be challenging due to the extensive range of variants tailored for different applications. However, as noted by J. C. Marè [89], several common elements are typically present in EHA systems for primary flight controls. These include

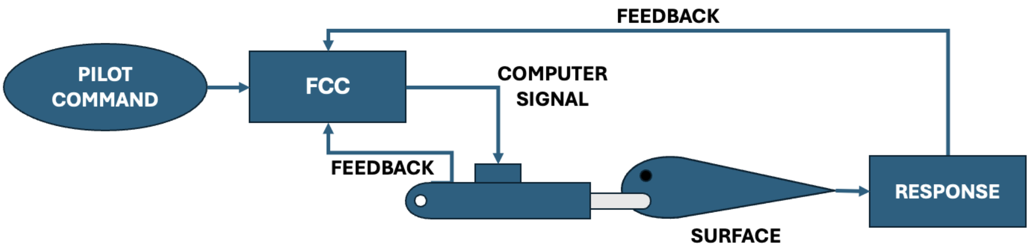

Fig. 1.3 A diagram of a FBW control loop.

- a servo valve that serves as the interface between hydraulic and electrical systems
- a main jack
- control electronics
- sensors and transducers

The servo valve itself is the core of the EHA and it is often selected between three main categories: Direct Drive Valves (DDVs), jet-pipe valves, or flapper-nozzle valves, each with its own advantages and disadvantages. Then, according to the aircraft category, mission, and size, the command from the pilot's control stick to the servo valve is transmitted through either a mechanical link or Fly-By-Wire (FBW) technology.

The mechanical link involves a complex arrangement of pulleys and rods that connect the control stick to the input interface of the servo valve. This cost-effective and time-proven method is employed for smaller control surfaces where the distance between the pilot and the aerodynamic surfaces is not too extended. However, for commercial aviation aircraft or for platforms where high performance is required, a mechanical transmission from the cockpit shows some limitations. In fact, higher aerodynamic loads on surfaces lead to a higher hinge moment, which can become unbearable for the pilot. Moreover, the responsiveness of the system can also suffer. As a result, the shift towards heavier and faster platforms in the '60 and '70 led to the introduction of FBW systems, where a low-power electronic signal generated by pilot input transducers in the cockpit controls the servo valve via the Flight Control Computer (FCC) [90, 91]. The FCC processes these signals and transmits them to the input connectors of the EHAs. Signals are brought to the main actuator block via cables, which are much lighter than rod and pulleys and require considerably less scheduled maintenance. The servo valve finally manages the hydraulic flow within one of the two chambers of the hydraulic cylinder, hence controlling the position of the jack [89, 92]. In fact, once received by the EHAs, the control signals are managed by control electronics that include logic systems typically composed of transducers (e.g. Linear or Rotary Variable Differential Transformers (LVDT/RVDT)) which provide feedback on jack position to close the position control loop often employed for primary flight controls. Additional internal control loops may monitor speed, spool position, or current levels.

If the control reaches the actuator servo valve via low power electronic signals, the hydraulic system provides hydraulic power and is connected with the servo valve through a complex network of secondary valves (such as bypass, isolation, and shut-off valves) along with additional components like accumulators and filters that ensure continuous hydraulic power while providing essential safety features. Moreover, some configurations incorporate additional sensors within the EHA body itself to monitor and/or log parameters such as pressure levels, debris presence, metal particle detection, and spool position. The information collected from these sensors could be logged into a HUMS, which can represent a solid base for the development of PHM solutions. A diagram of a FBW system is shown in Figure 1.3.

As mentioned above, the most commonly used types of EHAs include flapper-nozzle, jet-pipe, and Direct-Drive Valves (DDV) powered EHAs [93]. In particular, DDV controlled actuators are a specific type of actuators in which the spool is not moved due to differential pressures created upstream by a previous stage but directly by a Linear Force Motor (LFM) or solenoid [65]. This characteristic makes them conceptually different from flapper-nozzle and jet-pipe controlled EHAs. On a DDV, when an electrical command signal is applied, it generates a force that moves the spool to one side or the other, creating a pressure differential. This movement allows hydraulic fluid to flow to specific ports, which in turn controls the actuator's motion. In DDVs, the spools are engineered with even higher precision than other types of valves, ensuring that even small movements result in immediate changes in flow and actuator response. The presence of a LFM or solenoid eliminates the need for a pilot stage, thus improving response and pressure-independent dynamics. In some models, the spool is balanced by a balancing spring. A comprehensive discussion of EHA architectures and operating principles is provided by Marè [89] and is not repeated here for brevity.

1.10 Selected Case Study

This thesis approaches the complex and challenging task of developing a PHM/CBM framework tailored specifically for a case study in the aerospace field. The focus of this research is on the flight control actuators employed in an Advanced Jet Trainer (AJT) aircraft, a sophisticated platform built by Leonardo SpA designed to train military pilots in their advanced stages of training (Figure 1.4). Among the various components of this aircraft, the all-moving Horizontal Tail (HT) Electro-Hydraulic Actuator (EHA) is selected for the analysis.

The PHM framework developed in this work aims to predict Unscheduled Removals (URs) of the AJT HT EHA. This represents a clear distinction from much of the PHM literature, where the main objective is typically the prediction of specific actuator faults or failure modes. In the present industrial context, the primary interest is instead to anticipate those conditions that may lead the actuator to generate a no-go and therefore require its unscheduled removal, with direct consequences on asset availability and overall fleet management. The problem is therefore addressed at a higher system-management level, rather than at the level of isolated fault diagnosis. Moreover, as discussed in Sections 4.2.2 and 4.2.2, tracing each UR back to a specific fault mode or removal cause was not possible because of the limited statistical population, the presence of missing or incomplete information in the Equipment Registers, and the lack of actuator-level signals needed for a more detailed diagnosis. Consequently, the objective of the framework is to identify the likelihood of future UR events as a whole, rather than to classify the specific reasons that caused the unloading.

The selection of the HT EHA as a focal point for this thesis is driven by three main factors: its critical role in maintaining aircraft control during flight operations, its identification by Leonardo SpA as a potential pilot study equipment and the research gap related to PHM framework for EHAs encountered in the literature. In particular, this selection has been guided by in-house criticality and bad actor analyses among various aircraft components, emphasizing its importance in enhancing platform availability, reducing down-times, and streamlining maintenance processes.

Fig. 1.4 The AJT under study [1].

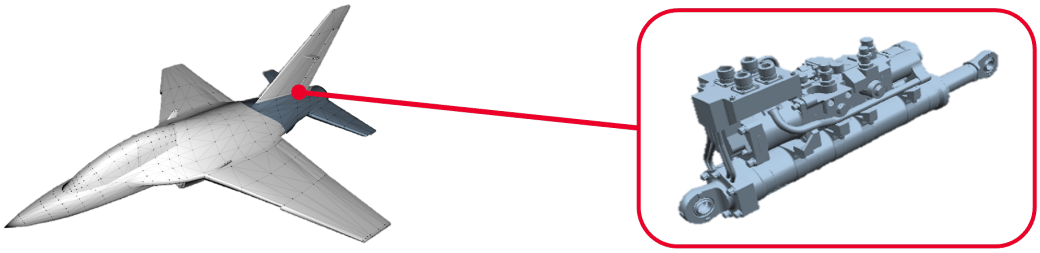

Fig. 1.5 The case study involves the EHA used to control the AJT Horizontal Tail (HT).

In addition, the literature panorama involving the development of PHM framework employing operational data for EHAs is definitely scattered and highlights a consistent research gap that requires attention. In fact, while diagnostic strategies for EHAs using system signals are well-documented [94], studies focusing on prognostics remain limited as already demonstrated and discussed in Section 1.8. This work seeks to bridge that gap by advancing innovative maintenance capabilities for EHAs.

The thesis will show how data collected from the aircraft can be integrated into PHM maintenance strategies of a specific component. Using historical performance data and real-time monitoring, the framework aims to enhance decision-making processes related to maintenance scheduling and asset management.

Furthermore, this research will address the very challenges associated with implementing PHM strategies in legacy systems. The ability to extract valuable insights from existing datasets, often collected for different aims, will be a significant aspect of this study. By focusing on the HT EHA, this thesis aims to contribute to the broader understanding of how advanced data analytics can optimize maintenance practices in aviation, ultimately leading to improved efficiency in flight operations.

The AJT unders study is a modern twin-jet-engine, tandem-seat lead-in fighter training platform that features advanced digital flight controls and avionics. The launch customer for the AJT was Italy, which placed an order in 2009, whereas currently, this AJT is already in service with several prominent air forces in the world. The development of this AJT was driven by a comprehensive analysis of the training requirements imposed by new combat aircraft. This analysis focused on factors such as maneuverability, thrust, human-machine interface, and avionics. As a result, the aircraft under study possesses a vast flight envelope, a high thrust-to-weight ratio, and exceptional maneuverability even at high angles of attack. This makes its handling characteristics optimal for training pilots whose final platform may be the Eurofighter Typhoon or the Lockheed Martin F-35. Consequently, the AJT under study is considered an ideal training platform to seamlessly transition to these advanced combat aircraft.

1.10.1 AJT FCS

The AJT FCS is tasked with aircraft maneuverability and controllability. The FCS operates as a full-authority, full-time, four-channel redundant digital fly-by-wire system. The main interactions between the four FCS and other systems and subsystems are shown in Figure 1.6. Starting from the top, the pilot commands on the pitch, roll, and yaw axes are captured through a set of Pilot Input Transducers (PIT) and passed to the FCS FCCs. Each FCC executes the aircraft control laws based on the received analog, digital, and discrete inputs, subsequently generating the required actuator commands for the control surfaces. These systems not only calculate the necessary actions for the actuators but also incorporate an actuator servo loop control, monitoring logic, warning and caution logic, Built-In Test (BIT) capabilities, input monitoring logic, and aircraft state sensor interface processing. Additionally, all operations conducted by an individual computer are subjected to cross-checking via Cross Channel Data Links (CCDLs). The command is then passed to the relevant actuation systems, which are powered by hydraulic and/or electric power provided by the hydraulic and/or electric system. The change in the position of the surfaces affects the flight dynamics. A loop back through the position feedback, an Attitude

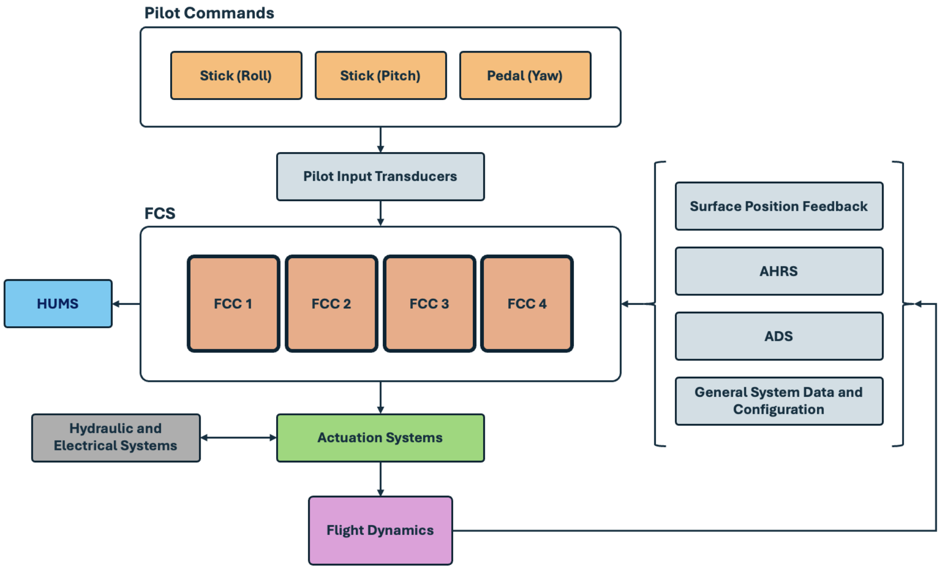

Fig. 1.6 AJT FCS architecture, elaborated starting from [2]

and Heading Reference System (AHRS), an Air Data System (ADS), and general system configuration data close the loop with the FCCs.

The AJT is equipped with a state-of-the-art Fly-By-Wire (FBW) hydraulically powered conventional suite of flight controls which provides excellent maneuverability and flight qualities. For the longitudinal plane, the AJT employs the all-moving HT and a set of LEFAS whose continuous movements enhance the longitudinal maneuverability. The all-moving HT design integrates the functions of both the stabilizer and the elevator into a single, cohesive moving surface, commonly referred to as a stabilator. This advanced configuration allows for enhanced aerodynamic performance and greater control. By merging these two critical components, the stabilator not only simplifies the design but also reduces the overall aerodynamic drag. Each stabilator is controlled by its own independent EHA, utilizing a tandem DDV controlled actuator. This actuator is integrated with the aircraft main hydraulic system via a dedicated control module. The design of this system ensures a high level of redundancy, which is critical for maintaining operational safety and reliability.

Modern advanced aircraft, such as the AJT, utilize parametric data acquisition and processing systems [22]. In the case of the AJT, this system is called Health and Usage Monitoring System (HUMS). These systems are designed to collect, store, and analyze flight data from various onboard sensors. Initially developed in the 1990s for rotary-wing aircraft, HUMS are closely aligned with PHM and CBM principles. Their primary role is to enable fault detection during or after flights by performing preliminary analyses on raw sensor data. In this regard, HUMS can be seen as a foundational PHM framework operating at a high level of abstraction.

HUMS are now standard across many aircraft platforms, with core functionalities shared despite slight variations in their implementation. For example, the Royal Netherlands Air Force (RNLAF) employs the F-16 Aircraft Condition Evaluation (FACE) system in its F-16 fleet [86]. However, legacy systems often feature data logging modules tailored to specific applications, which may not meet modern standards. The AJT HUMS, for instance, was specifically designed to support Structural Health Monitoring (SHM) and fatigue tracking [95–97].

Transitioning from HUMS to a fully integrated PHM system is not without challenges. Common limitations include inconsistent or low sampling rates, missing data, and limited sensor coverage—problems also noted by Kappas and Frith in their analysis of HUMS systems [25]. These challenges highlight the complexities involved in adapting existing systems for advanced PHM applications. This thesis aims to address these challenges by developing a PHM system that leverages operational data to predict equipment health.

Chapter 2

Project Rationale and Workflow

2.1 Rationale for the Selected PHM Strategy

The proposed PHM system is distinguished by several key properties: it is data-driven, scalable, robust, grounded, and transferable. Furthermore, it utilizes aircraft-level data and is designed to be "as simple as possible and as complex as necessary".

Developing a PHM system from scratch is a very challenging project, let alone if the selected equipment is a legacy system loaded onto an operational platform [98]. To establish a robust, reliable and scalable system, a fully data-driven approach has been selected and adopted. Given the specific research questions addressed in this thesis, and without loss of generality, this decision is influenced by various factors, including the availability of design data, data quality, access to operational flight records, and the commitment to design the PHM system grounded on operational data. Additional complexity is added only where and whether necessary. This aligns with the adherence to the "Less is more" philosophy.

Many literature records show that, more often than not, a hybrid PHM approach is the winning one, combining the customized data driven part with the generalization ability of the model base section. This consideration is indeed true and is backed up by several studies. However, a "one size fits all" solution is a privilege that the PHM research area cannot afford to adopt. Developing a model for the EHA itself could have been possible, but the low and variable data sampling rate and the absence of the input command of the actuator would have made the practical integration of the model in the PHM scheme hardly applicable. Moreover, model design and validation data are lacking. Hybrid PHM is indeed the best solution, but only when it is possible and affordable to build a model.

The fully data driven approach has been conceived to work on the data available from a production aircraft in the way that this system could be switched onto every aircraft of every fleet without any change or retrofit. In the era of AI and machine learning (ML), a often sought-after property of a data-based methodology is the scalability and transferability of the solution. One of the benefits of the provided solution is the potential transferability to other equipment, systems, or even platform logging data in time series format. The selected approach is indeed an across system solution to an equipment problem. The added value of this decision must not be overlooked: the potential to bring assets on the PHM journey only with the integration of existing data represents a significant enhancement to the project. Many PHM works require the implementation of additional sensors on the asset being monitored or a change in some operative procedure. We chose to stay grounded and propose a scalable and modular solution that could be easily integrated with other fleets and platforms.

Often, equipment level PHM systems are built on equipment level data. This consideration may seem obvious, but it masks the idea that, in order to track down the health status of a system, there is the necessity of monitoring the equipment first. In laboratory conditions or in modeling framework this is easy, but working on legacy equipment means that "you have to do the best with what you have". In the case of this thesis, no equipment level signal was available except from the actuator position, obtained by the FCC and logged in the on-board HUMS system. This consideration opens up a completely new set of challenges since, apart from the actuator position, all the other signals are aircraft level signals. In this thesis, aircraft-level signals are measurements available at the aircraft platform level (i.e., recorded by aircraft-wide monitoring/avionics systems) and not tied to a specific subsystem or component. In other words, the project involves tracking down equipment level health status leveraging aircraft level data.

Finally, as already stated, this framework targets the prediction of URs of the AJT HT EHA, rather than the diagnosis of specific actuator fault modes or failure mechanisms. The objective is to anticipate those conditions that may lead to a no-go and consequently to an unscheduled unloading, since these events have a direct impact on aircraft availability and fleet management. The classification of the exact underlying removal cause could not addressed, because the available data do not support a reliable fault-level analysis: the statistical population is limited, the Equipment Registers contain missing or incomplete information, and the actuator-level signals required to trace each UR back to a specific fault mode are not available. Further details on the fault modes are addressed in Section 4.2.2.

Building upon these insights, the next chapter will expand these principles into a robust and modular strategy that could handle sparse industrial data, such as those in the available dataset.

2.2 Strategy Overview

This section offers a concise overview of the proposed strategy, which will be elaborated upon in greater detail in the subsequent section. Building on the previous discussion of the rationale of the design, the selected strategy employs the first four statistical moments—mean, variance, skewness, and kurtosis—to analyze the asset signals collected after each flight. The statistical moments summarize the behavior of the signals during each flight, enabling a monitoring of asset performance over time. This approach condenses the behavior of each flight into a set of manageable parameters, serving as a data concentration technique. Moreover, the system needs to be robust to handle relatively low data quality and to exploit maximum information from sparse data.

The trends of these statistical moments are then transformed into cumulative features that represent the usage of the equipment. These cumulative features are ranked based on their correlation with URs by creating a probability distribution of the magnitude increase of cumulative features between two URs. This ranking serves two purposes: it identifies the most informative cumulative features in a clear and traceable manner and establishes a threshold distribution for future predictions. Since the distribution of magnitude changes between URs is now known, it becomes possible to compare trends with the threshold distribution to determine if the predicted values fall within that range and assign an UR risk.

Once the most influential cumulative features are identified, they can be projected into the future using Monte Carlo simulations. These simulations rely on a sequence of Clustered Mission Types (CMTs) identified through clustering analysis of altitude records from a subset of flights, which revealed 100 clusters of similar flights. Understanding this cluster distribution allows one to predict future flight sequences by selecting clusters from this distribution in an informed way. Nevertheless, it allows to predict the future equipment usage for a specific set of mission by setting customized CMT future sequences.

After projecting the cumulative features into the future with associated probabilities, their values can be compared against the previously established threshold probabilities. Custom formulas can then be used to calculate the Probability of Failure (PoF) and assess risk trends by comparing these probabilities with the threshold distribution. The resulting risk trend, which ranges from 0 to 100, provides actionable insights that can enhance decision-making and improve fleet operations from a short- to mid- and long-term perspective.

2.3 Workflow

The adopted workflow has been inspired by the standards and practices reported in Section 1.7, influenced by similar literature studies, and refined to match the specific characteristics of the case study.

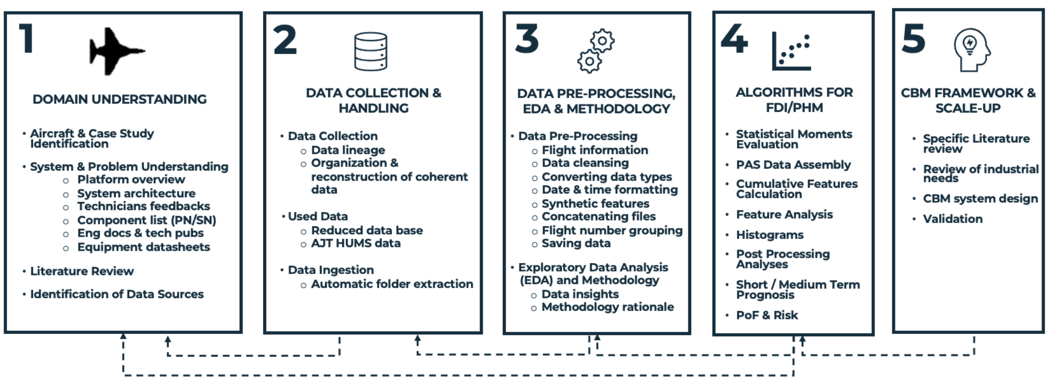

Fig. 2.1 Project workflow, ad inspired by literature similar PHM projects and experiences

The research work has been divided into 5 main steps. The first step marked the knowledge acquisition phase. This step is crucial as it involved the choice of the aircraft and case study and the overall understanding of the domain. The case study, the AJT HT EHA, has been selected by Leonardo SpA considering the improvement of the Mean Time Between Unscheduled Removal (MTBUR) of the equipment and the scientific relevance of the topic, as already exposed in Section 1.10. In PHM projects, it is customary to divide the competences of the SME and the data scientist. However, in the context of this project, the two roles are intertwined. That is why a central step of the project involved getting acquainted with the platform and the system under study. This involved daily interactions with contact points in industry units. Operative support, logistic engineering, SHM, FCS, and quality units have contributed significantly to obtain useful information and documentation. After the identification of the Part Number (PN) and Serial Number (SN) list, in this phase, technical documentation and publication involving the platform and the system have been analyzed to conceive a possible strategy and trying to understand the possible path. A particularly relevant step has been the feedback received from maintenance technicians and ground personnel working every day with the aircraft. A specific literature review on the topic has been carried out and then constantly updated throughout the three years to establish the state-of-the-art and catalog the records [3]. The review results have already been reported in Section 1.8 and underscore a scattered and scarce literature base with mostly component level contributions mainly leveraging laboratory or synthetically simulated data. This first phase helped to better identify possible solutions and strategies, as well as to understand data sources and operative procedures. The identification of data sources required an extensive effort in the direction of finding the possible data in very unit and database and concludes this first step.

The second phase regarded data collection and handling. After having identified all useful data sources in the first step, in this phase, data had to be gathered, collected, and organized. In addition, the data lineage has been reconstructed for the specific PN under investigation. Data lineage can be defined as the detailed tracking of data journey throughout their life-cycle, with the aim of understanding and monitoring data flows from origin to final destination. By mapping how data are created, processed, and used, it is possible to highlight any possible data hole or missing sources to provide better insight. The data sources identified in step 1 included a wide variety of formats, types, and quantity of data. More often than not, data quality can be an issue when dealing with operational real-life data as data can be corrupted, filled with data recording errors, saved in non coherent formats, etc. Data are transcribed by hand or the data transfer is carried out offline and in a non synchronized way. Furthermore, data involving the same asset but saved in different databases may be incoherent due to data saving settings. This is why reconstruction of a coherent data flow has been pivotal in driving the next phases. After having gathered all data sources in a single data repository, a reduced data base was created with data that was deemed to be the most useful one. The "golden source" is definitely represented by the flight records logged in the aircraft HUMS. Step 2 has been completed with the data ingestion step in which a routine was set up to automatically extract HUMS data from the folders.

This leads to the third step, where data science becomes a priority. In this phase, Data Pre-Processing took place, along with an Exploratory Data Analysis (EDA) and the envision of a data mining methodology. Raw data has been cleaned to remove spurious information. Data types have been converted, date and time information has been formatted following the UNIX convention. Some additional features were created by combining the existing one, files have been concatenated, grouped on a flight number basis and then saved. An EDA revealed important insights on data quality, consistency, and sampling and was important to inform the next steps.

Step 4 has focused on the main data mining techniques which will be better explained in Chapter 5 and involved the calculation of statistical moments, the creation of cumulative features, related histograms and the calculation of the Probability of Failure (PoF) and risk trends.

Finally, risk trends were analyzed from a CBM perspective after conducting a literature review and analyzing the industrial need related to the scale-up of the CBM.

The workflow just explained is reported in Figure 2.1, which shows the sequence of steps carried out in the three year period.

The dotted arrows shown behind each block indicate that the development of a PHM framework, particularly when dealing with a legacy product, is not a linear process but rather an iterative one. For example, if it becomes apparent during Phase 4 that a previously unconsidered signal could be beneficial for analysis, a feedback loop is established to retrieve the necessary information from the data pre processing phase. Additionally, if further information is required, the PHM designer is encouraged to revisit earlier stages to acquire the specific knowledge needed, which can then be incorporated into subsequent phases.

These feedback loops should not be viewed as errors or setbacks, but rather as vital components of a continuous improvement process that is essential for the advancement of a forward-looking framework. Throughout the research project, feedback from technicians and experts has often inspired reflections leading to significant enhancements in the system, but only after thoroughly passing through the required processes again and again.

Chapter 3

PHM Functional Architecture

The developed PHM strategy is better visualized and clearly explained if a modular architecture divided into functional layers is adopted, such as the one presented in Figure 3.1.

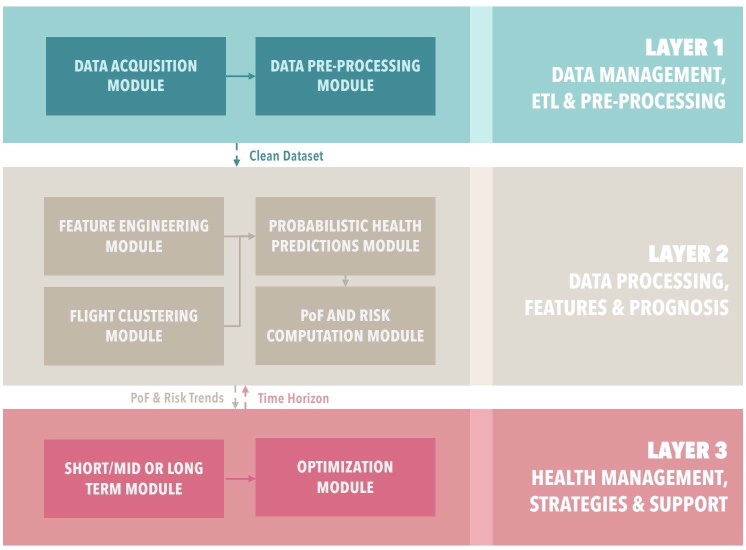

Fig. 3.1 PHM framework divided into functional layers and working modules

Dividing a PHM, or any software application, into layers is an optimal solution to visualize the objective of each layer and modules. Moreover, many standards concerning CBM, IoT, cyber-physical systems, and information technology applications in systems engineering and aerospace propose the organization of the PHM system architecture into integrated and modular layers. These standards have already been analyzed in Section 1.7. The MIMOSA group emphasizes the feasibility of a standardized layered architecture specifically for PHM systems in both the OSA-CBM and OSA-EAI. Additionally, SAE ARP 6883 (Guidelines for Writing Integrated Vehicle Health Management (IVHM) Requirements for Aerospace Systems), while not explicitly defining software layers, offers a structured methodology for IVHM system development that can be merged into a layered system. Finally, many literature records adopt a layered framework. As a result, inspired by literature studies and standards, the proposed architecture is conceptualized as a three-layered modular framework.

Each layer and module will be analyzed and presented in detail in the next chapters (one for each layer) and sections (one for each module). A short description of the overall framework is here reported for reasons of clarity.

- The first layer encompasses Data Management, ETL, & Pre-Processing
- The second layer is responsible for Data Processing, Features, & Prognosis
- The third layer involves Health Management, Strategies, & Support.

3.1 Layered Architecture

3.1.1 First layer: Data Management, ETL, & Pre-Processing

The first layer is responsible for the Data Management, Extraction, Transformation & Load (ETL) and Pre-Processing. It contains two working modules, the Data Acquisition module and the Data Pre-Processing module. Once the raw data repository has been acquired, pre-processed, and formatted, the clean dataset is passed to the second layer. The Data Acquisition module is tasked with overseeing the data ingestion process, including Extraction, Transformation, and Load (ETL), as well as the automatic integration of data from various sources and databases. It also manages data formatting processes. Once the data are formatted, they are transferred to the Data Pre-Processing Module, where flight information is extracted. This module is responsible for data cleansing and structuring, converting data types into coherent formats, and standardizing date and time information to a Unix-based system. Additionally, it concatenates various files and saves the processed data in both a Python dictionary of data frames and ".csv" format to support future inspections and analyses. During this stage, an Exploratory Data Analysis (EDA) is conducted to identify valuable insights within the data and to highlight any potential data quality issues.

3.1.2 Second Layer: Data Processing, Features, & Prognosis

Focusing on Data Processing, Feature & Prognosis, the second functional layer is the most complex, and it contains a total of 4 interconnected working modules, providing the PoF and Risk Trends to the next layer. This layer develops advanced probabilistic data models to derive and enhance value from data. It evaluates the current health status of the system and forecasts future behavior based on the current assessment. The Feature Engineering Module performs two tasks. On the first hand, it analyses the clean dataset and extracts a ranking of the most informative features obtained through the data manipulation of cleaned signals. On the second hand, it establishes a threshold distribution. In parallel, the Flight Clustering Module operates on three main functions. Firstly, it receives the clean dataset and operates a specific clustering routine to identify and cluster similar flight based on the flight path in Clustered Mission Types (CMTs). Moreover, it provides a sequence of potential future CMTs, which are then employed to do forecasting. Finally, it assigns an increment of each feature linked to each CMTs. All this information, along with the features obtained in the Feature Engineering Module, are passed to the Probabilistic Health Predictions Module, which is responsible for the feature probabilistic forecasting in the next sequence of flights with a guidance provided by CMTs leveraging Monte-Carlo simulations. After that, the probabilistic trends are compared with threshold distributions, obtaining PoF and risk trends [99, 100]. These trends represent the input for the third layer, called Health Management, Strategies, and Support, which exploits the obtained risk curve to enable short- or mid-term fleet optimization.

3.1.3 Third Layer: Health Management, Strategies, & Support

The third layer is responsible for integrating risk trends data into an existing industrial logistics and maintenance framework. The customizable design of the Probabilistic Health Prediction Module facilitates the seamless application of the PHM framework in various contexts, thus enabling tailored solutions that meet specific customer needs. The utilization of this information is closely aligned with the operational requirements and procedures of the end operator. Risk assessments and strategies contribute to shaping availability predictions for both short-term and long-term scenarios. By simulating a set of predefined CMTs in the short-term, it is possible to simulate the increasing UR risk informed by the operational data-set. This option gives the potential to foresee specific mission sequences and evaluate the impact of the CMT sequence on a specific EHA loaded on a specific aircraft. The consequences of these detailed analysis could then be integrated into a decision making approach providing tailored advisory on which aircraft is the best suited for the set of mission considering the risk of EHA URs. Additionally, a long-term perspective can be established by implementing “open-loop” CMT sequences, with their distributions informed by and drawn from historical statistical data. This approach allows for the assessment and estimation of the impact of long-term aircraft usage with positive outcomes on supply chain planning, warehouse management, etc.

A detailed architecture overview of the PHM framework is reported in Figure 3.2. Each functional layer and working module is further expanded to reveal the process block and the relationships between them. In the next chapters the architecture will be further analyzed and discussed.

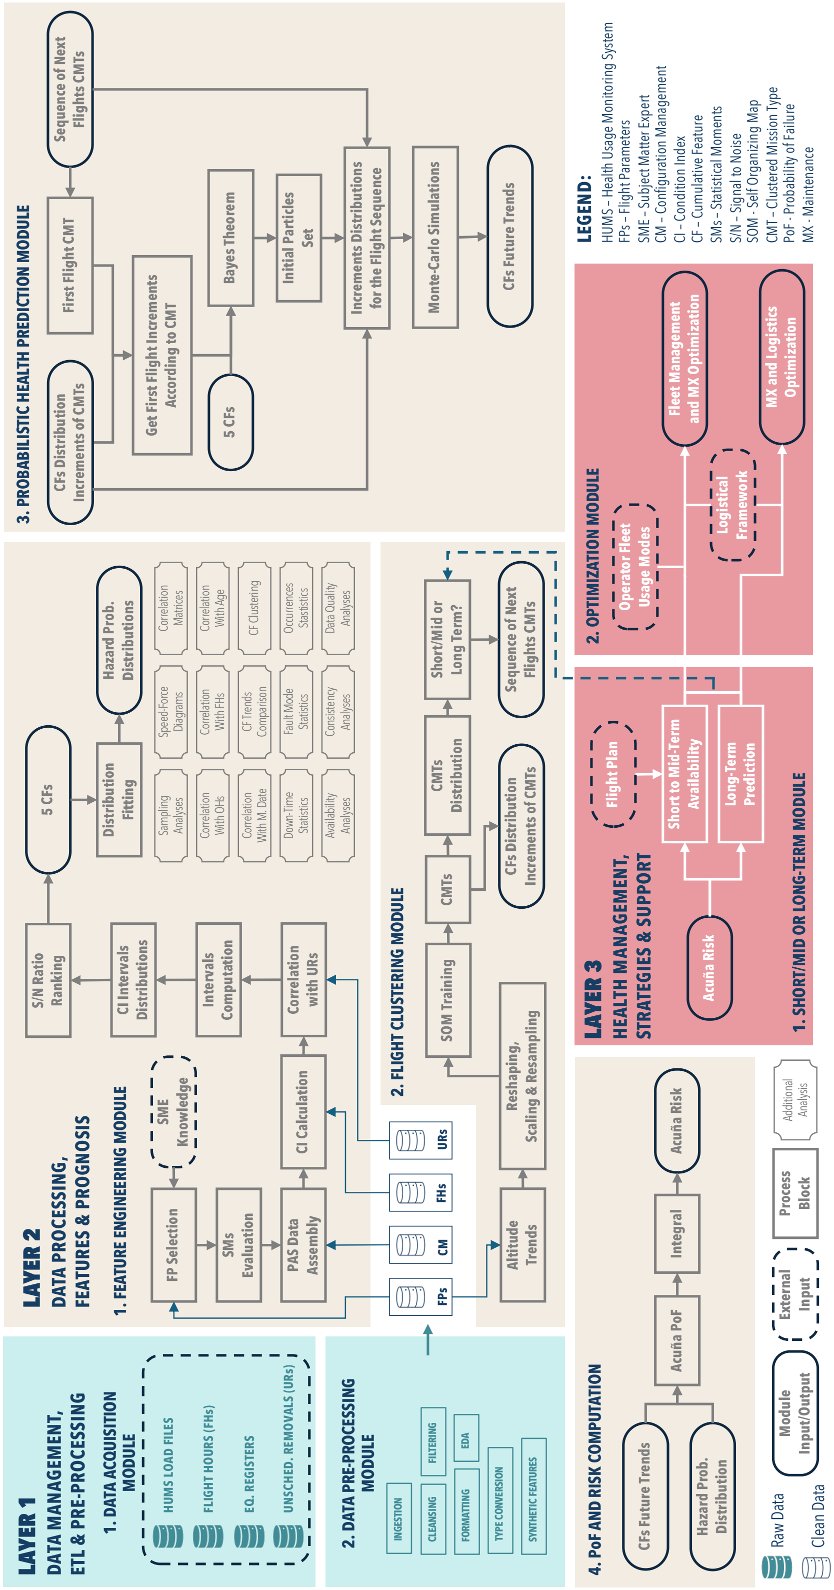

Fig. 3.2 PHM framework detailed architecture

Chapter 4

Layer 1: Data Management, ETL and Pre-Processing

4.1 Data Acquisition Module

The effectiveness of a PHM framework is directly dependent on the quality, quantity, and representativeness of the underlying data. Data serves as a crucial component in the development of any PHM or CBM program; therefore, careful attention must be paid to identifying data sources and collecting data.

This thesis focuses on the development of a PHM framework for legacy equipment. As a consequence, the data acquisition process is approached from the perspective of working with an existing legacy product. In this case, data acquisition and management focuses on selecting, identifying, and leveraging data that are already being recorded with legacy sensors. Designing a PHM framework for a new asset presents a completely different scenario, as the PHM designer would need to write and comply with specific requirements for the PHM system. Additionally, there would be opportunities to select specific sensors networks [101, 102], choose sampling frequencies according to specific sensitivity analysis [103], and collaborate with other design units to integrate the PHM system into the overall vehicle design in the best way possible. However, this is not the case for this thesis, as this project will limit itself to leveraging existing data to design the PHM framework.

The term "Data Management" effectively encompasses a range of activities performed on data before actual extraction and use. Activities such as reconstructing data lineage, acquiring data, storing data, and organizing stored data fall under the umbrella of data management. This is not an engineering activity strictly speaking, it is more aligned with management work, involving communication with individuals and reviewing technical publications in order to identify similarities between data sources utilized in other studies and the data available in the specific case study. It is important to note that the names and definitions of terms can vary significantly between organizations and institutions, which can lead to confusion that should be addressed at the beginning of the research project.

The data acquisition process required over a year and a half of work within the three-year research project. Although this duration may appear exaggerated, the literature indicates that it is actually on the shorter side, as data acquisition and management in PHM projects can take up to 90 per cent of the total program time [104]. After this phase is completed and all possible data sources have been identified, the data must be organized.

A common division of data sources distinguishes between operational data sources and non-operating data sources. Operational data involves all those data sources that provide data linked to the effective use of the asset (e.g. maintenance records, flight data, etc). On the other hand, non operational data involve all those data sources which are related to the equipment under investigation but are not related to the use of the equipment (e.g. technical publications, reports, etc).

Operational data can be further categorized into event data and condition monitoring data. Event data pertain to occurrences that affect the asset, such as maintenance logs and unplanned repairs, while condition monitoring data involve flight data and measurements that track the health condition or status of the physical asset [105].

Component-related maintenance events can also be divided into two categories: failure events and censored events [106]. Failure events arise from unexpected component failures and need unscheduled corrective maintenance, while censored events involve the preventive replacement of a component at predetermined time intervals.

Technical publications, technical drawings, equipment data sheets, and other documents detailing the overall architecture of the system are critical to understanding the principles of the platform and the interactions between subsystems. Additionally, operational data, including flight data and maintenance reports, constitute the primary dataset necessary to develop algorithms and methodologies.

In accordance with the previously mentioned standards and the literature records, a preliminary list of the data required to effectively develop a PHM/CBM flow chart for the case study has been compiled. Consequently, a comprehensive review of all relevant engineering and operational documents is essential to identify potential data sources. To identify all data sources and the several interconnections between them, a data organization diagram has been produced (Figure 4.1) and extensively explained in a published conference paper [5].

An approach to interpret the diagram presented in Figure 4.1 is to follow the blue boxes that delineate the key steps. The initial phase focuses on system and problem understanding and then progresses to the definition of the PHM/CBM system and the potential development of a model. Additionally, supplementary data sources can be utilized to enhance the research project with further information. The next paragraphs present in details the principles on which the diagram was designed.

Starting from the top center, the root box is represented by the PN and SN for the selected subsystem and aircraft. They have been essential in defining the system architecture and, as such, they are linked to a wide range of data boxes, such as Safety and Reliability (S&R) analyses, all system level information (control laws, hydraulic system details, data exchange protocol control architecture, structural data, item removal causes and data, etc.). Safety and Reliability (S&R) analyses, along with Failure Modes Effects and Criticality Analysis (FMEAs) and Fault Tree Analysis (FTAs), have played a key role in identifying critical failure modes and assessing the severity of the issues being investigated.

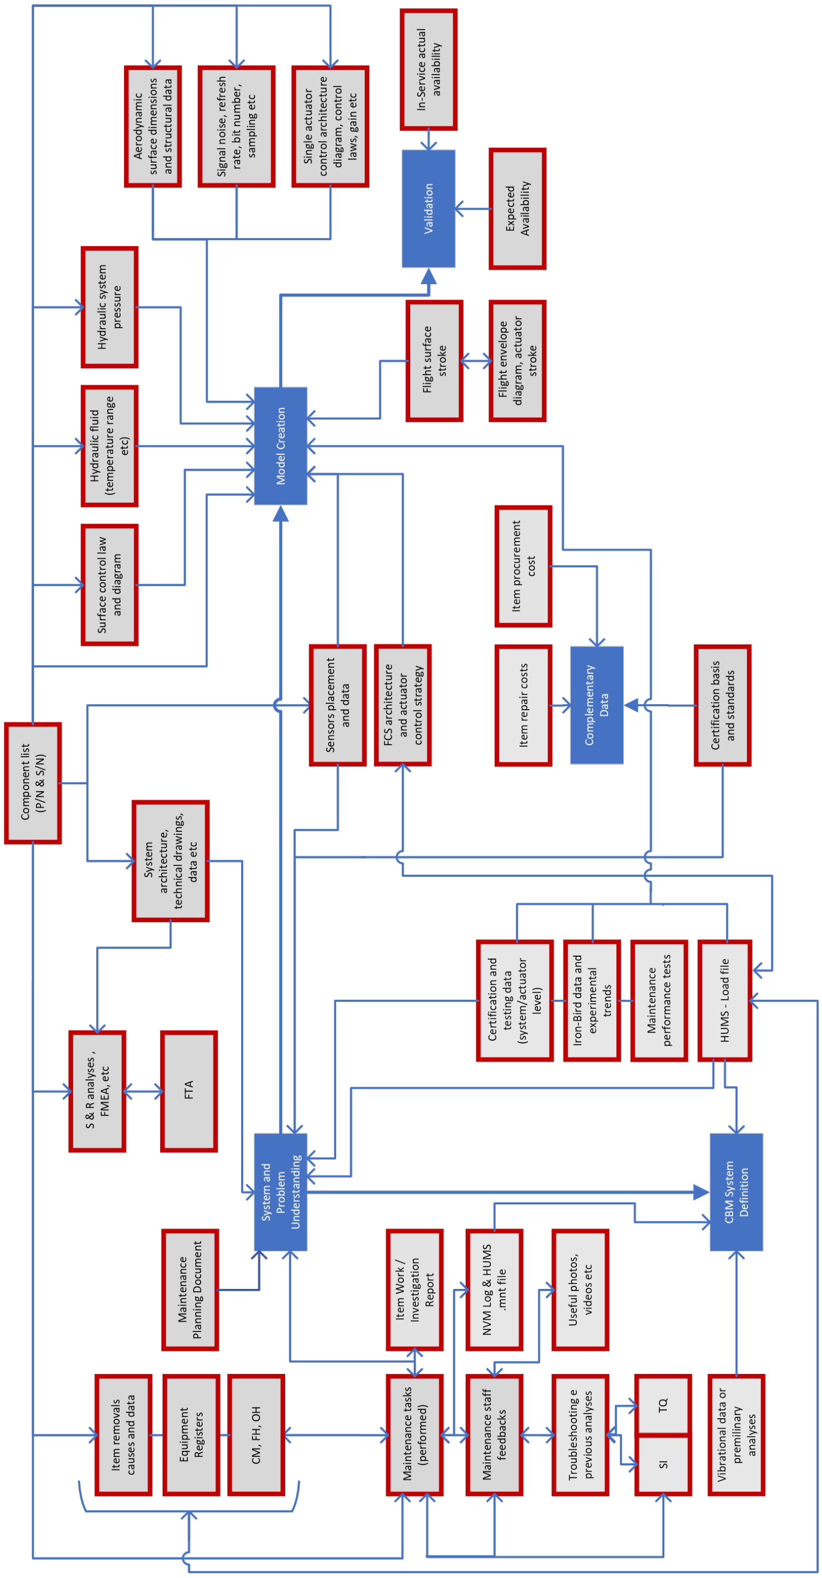

Fig. 4.1 Data organization overview

Following the central arrow, the section titled "System Architecture, Technical Drawings, Data, etc." includes a variety of engineering documents that are crucial for understanding the system and the challenges faced, especially while integrated with information from S&R analyses. At the same time, moving to the left part of the diagram, information on component replacement reasons has been consolidated with operational Equipment Registers, Configuration Management (CM) data, as well as Flight Hours (FH) and Operating Hours (OH) data. Going lower, this information is linked to Maintenance tasks, then to staff feedback, investigation reports, and finally to the main data sources: Non Volatile Memory (NVM) and HUMS files. This cluster of boxes, shown on the upper left corner of the diagram, represents the first part of core data base of the project. From that group of boxes, we were able to obtain information on troubleshooting and previous analysis, photos, and videos. This information has been combined with vibrational data and preliminary analyses, including Segnalazione Inconvenienti (SI), Occurrence Reporting, and Technical Queries (TQ). SIs and TQs provide formal channels for customers to request additional support from the engineering division.

The HUMS .mnt files document system-level faults and alerts that occur during flights. NVM logs capture FCC fault codes and are downloaded from the aircraft under specific circumstances; when available, they provide further information on reported faults alongside the associated .mnt files. Maintenance performance tests and Iron Bird test data have also been aggregated with certification test data. All these mentioned data sources lead to the blue box "System and Problem understanding", after which it has been possible to formalize the "CBM System Definition" where the previous blue box information converges with other already performed preliminary analyses and the central low cluster of data sources (representing the second main data source cluster). In this group there are certification data, Iron-bird data, maintenance performance tests, and HUMS Load file, which contains time series data. A link is placed between the HUMS Load file and the CM, FH, and OH box because these information must be considered together.

Moving to the right-hand section of the diagram, certain data sources could potentially be used as model parameters for high-fidelity or low-fidelity physical models. These include sensor placement, surface control laws, FCC algorithms, hydraulic fluid properties, aerodynamic surface dimensions, structural data, Flight Control System (FCS) architecture, control architecture diagrams, flight surface details, actuator stroke information, and flight envelope diagrams. Availability data

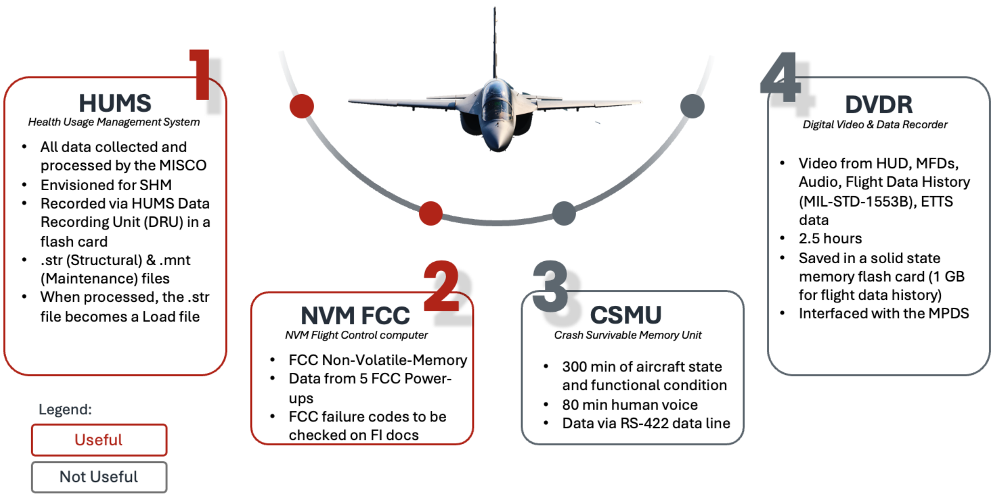

Fig. 4.2 View of potential in-flight data sources

can help validate the PHM system. Additional supplementary data includes item cost information, certification basis, and compliance standards.

Other operational data sources typically used in the development of PHM strategies for current and legacy equipment may include the Crash Survival Memory Unit (CSMU) or the Digital Video and Data Recorder (DVDR) (Figure 4.2). However, these sources were excluded from this study due to inconsistent data retrieval methods that are triggered only by specific events rather than through regular and established processes.

This organization of data has been crucial in clarifying data requests within the industry and visualizing interactions within a complex system and operational environment.

Against this backdrop, the organization of the identified data repository is also reported in Figure 4.3, where the data sources are grouped into macro-categories and divided into Operational Data and Non Operational data. To all intents and purposes, the dataset is a multimodal dataset in the sense that it comprises information from various modalities that collectively describe features of the same function. In fact, multimodality, in the context of information and data, refers to the presence of multiple modalities within a single dataset [107]. A hot research topic in the PHM community is related to multimodal learning. Multimodal learning can be defined as the process of deriving valuable insights from data that encompass multiple modalities, while also taking into account the interactions between these modalities. Engaging with various modalities is essential, as real-world information frequently involves multiple forms of data. In this sense, a PHM data can include multimodal sources such as time-series signals, maintenance reports, images taken from technicians, performance tests performed during scheduled maintenance routines, etc. In the case of the AJT datasets, the various modes can be, however, reconducted to an "oligo-modality" dataset as the main data source is represented by HUMS files, which contain signal time series. Maintenance reports are quite scarce, do not provide important information on the reason for equipment unloading, and have only been used to statistically characterize the most frequent failure modes. In the dataset, there are photos taken by maintenance technicians, but they are not directly linked to the actuator per se, as they are related to the HT. Performance tests are often not kept by the maintainers and, therefore, there is no historical basis to work on multimodal learning. Nevertheless, in this thesis, the inherent principle of multimodality is embraced as much as possible by integrating equipment registers with time series. If in the future a deeper and more multifaceted dataset will be available for this case study, a new multimodal strategy could be approached by keeping the same modular framework.

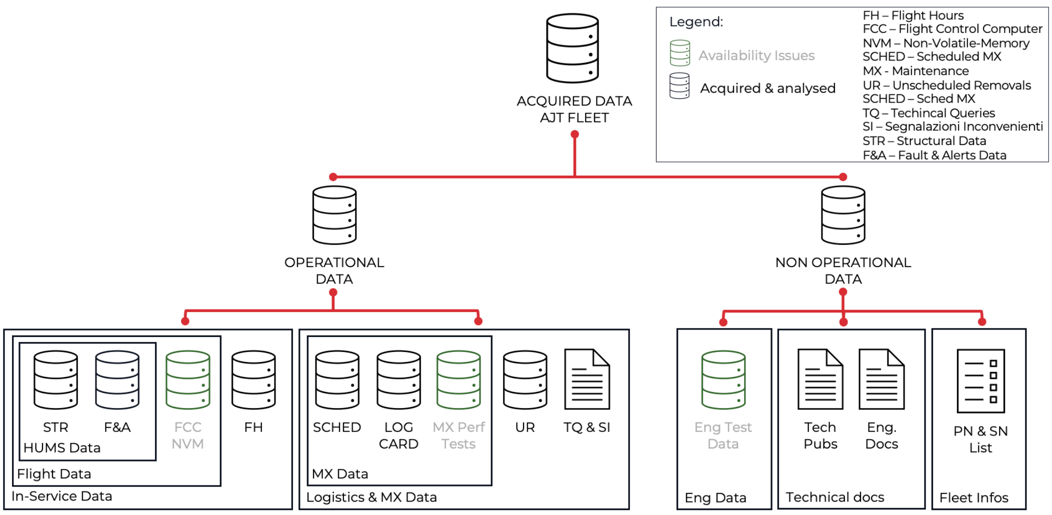

Fig. 4.3 Data repository: each source has been grouped up and divided into operational data and non operational data

HUMS data encompasses both STR and F&A files, while Flight Data includes HUMS data and FCC NVM. The In-Service Data include Flight Data and the

Flight Hours register. Conversely, Logistics and Maintenance Data consist of the Maintenance data category, which includes the Scheduled Maintenance register, Log Card register, and performance tests conducted during maintenance actions, along with the UR register and TQ & SI. Engineering data, technical documentation, and fleet information are classified as non-operational data.

Following a comprehensive review of each record and source type, a pruning operation was conducted to keep only the most relevant sources for informing subsequent steps. For example, the FH register was deemed superfluous, as the duration of each flight can be readily obtained from the HUMS file by comparing flight start and end times. Additionally, the FCC NVM has been found to be difficult to obtain and provides limited additional information. Similar considerations apply to maintenance performance tests and engineering test data. As a result, the diagram shown in Figure 4.4 shows the selected reduced data sources.

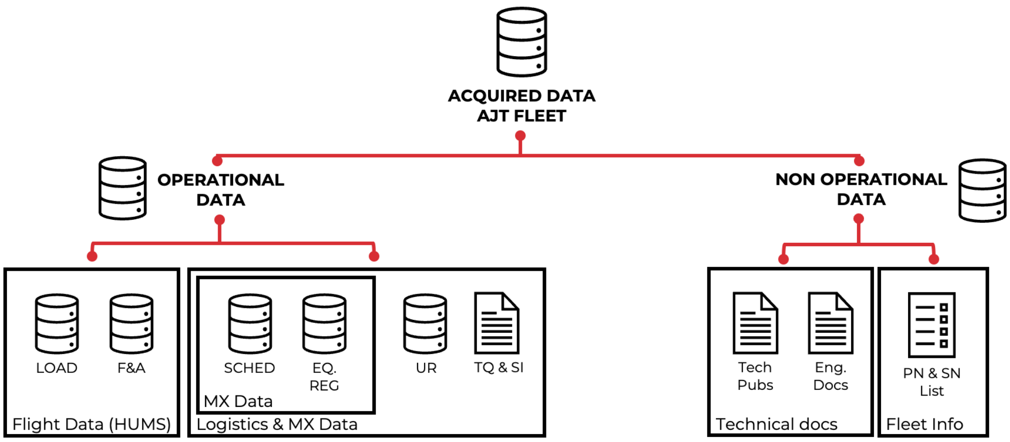

Fig. 4.4 Reduced data set with only relevant data sources

In the case of the AJT EHA, raw data is available in various formats. HUMS LOAD (also named STR - for structural) and F&A files are provided in TXT format, while Scheduled Maintenance registers are available in both PDF and Excel formats. Equipment Registers, as well as URs, are only available in PDF format. The TQ and SI documents are also formatted as PDF, and technical documents are provided in PDF format, with Fleet Information available in Excel format.

To all intents and purposes, the magnitude of the data to be analyzed leads to the notion of Big Data [108, 109]. Big data is often defined as "extensive" or "vast" databases. However, this definition is not enough to define what is big data and what is not. For instance, a more complete definition may involve the concept of computational power needed to store and process it. Furthermore, if a dataset can be seen as big data, it is determined not only by its size but also by its content and intended applications. A common definition of what big data is involves the "3 V's": Volume, Variety, and Velocity. Volume refers to the mere amount of data that can significantly impact processing capabilities. Variety is related to the diversity of data types and structures within a dataset. Lastly, Velocity refers to the speed at which data are generated and must be processed. This project entire data repository can be defined as a big data database since the data volume is definitely extended, spanning more than 10 years of data from a fleet of 22 aircraft for a total of more than 25000 flight Hours. Variety is definitely a characteristic of the data repository since time series (categorical and numerical) data have to be integrated with written documentations, maintenance reports, etc. Velocity is the least pronounced feature in the database because the acquisition frequency is not too high by today's standards. In any case, each file for each flight contains a considerable amount of points.

The data analysis framework is based on Python 3.8.18, using VSCode. The most commonly used packages have been installed and used along with specific packages to perform custom made operations.

4.1.1 HUMS LOAD File

An insight is required for the HUMS LOAD file, which is the main operational data source. The LOAD file is the result of a decoding and transformation process (from binary to exadecimal) starting from a coded file, which is saved in a memory card from the aircraft Data Recording Unit (DRU) along with the F&A file. These files can be downloaded after each flight and then analyzed accordingly in the HUMS ground segment, constituted by a Ground Support System (GSS). In particular, after every flight, the F&A file is analyzed to see if there were any alerts notified during the flight. If this is the case, then the specific system is analyzed more in detail and by looking at the trends reported in the LOAD file if need be. In the case of the FCS, there is also the possibility of accessing the FCC NVM where the FCC logs are saved for a limited number of flights. However, it should be noted that this procedure is performed only in the event of a reported alert in the F&A file and not regularly. Data contained in the LOAD file, once decoded, are in the format of txt files where all time-series data is saved. Moreover, some categorical data is also saved. (e.g., Weight-On-Wheels (WoW) or additional tanks on two or more aircraft pylons, etc.). In order to optimize data saving space, as often carried out in these conditions, a data compression algorithm is applied during the saving process by the HUMS airborne segment. It has to be noted that, as already explained in Section 1.10, the AJT HUMS was designed for SHM, for which the dynamics are characterized by much lower frequencies.

In this specific case, a Peak-Valley (PV) approach had been selected by AJT designers in the '90s. Under this acquisition logic, the system does not record signals continuously during the flight. Instead, each monitored parameter is compared in real time with predefined trigger thresholds, and data are stored only when at least one signal in the selected subset (e.g. Equivalent Airspeed, Mach Number, Angle or attack/sideslip, angular speeds and accelerations, etc.) shows a sufficiently large variation to activate the trigger. After a recording event, a dead band is applied, meaning that no new sample is saved until the signal changes by an additional prescribed amount. In practice, the process is therefore based on three steps: continuous monitoring of the selected parameters, storage of data when a threshold exceedance occurs, and temporary inhibition of further recordings until the dead-band condition is overcome. This strategy allows the recorder to capture the most relevant extrema and transitions of the flight profile while avoiding redundant storage of very similar points and limiting memory usage.

An example of the PV strategy is shown in Figure 4.5. The upper curve represents the original time history, which is continuously monitored, whereas the lower curve represents the compressed signal actually stored by the acquisition system. A new point is recorded only when the signal variation exceeds the dead-band threshold with respect to the last stored value; otherwise, intermediate fluctuations are ignored. In this way, the PV logic preserves the main peaks, valleys, and significant transitions of the original signal while discarding small variations, thereby reducing the amount of stored data. The loss information can be significant as there may be intervals during which no exceedance is recorded. Conversely, there may also be instances where the DRU receives continuous data due to persistent exceedance triggering. This highlights the high variability in the recording frequency that will be the subject of further analysis in Section 5.1.1.

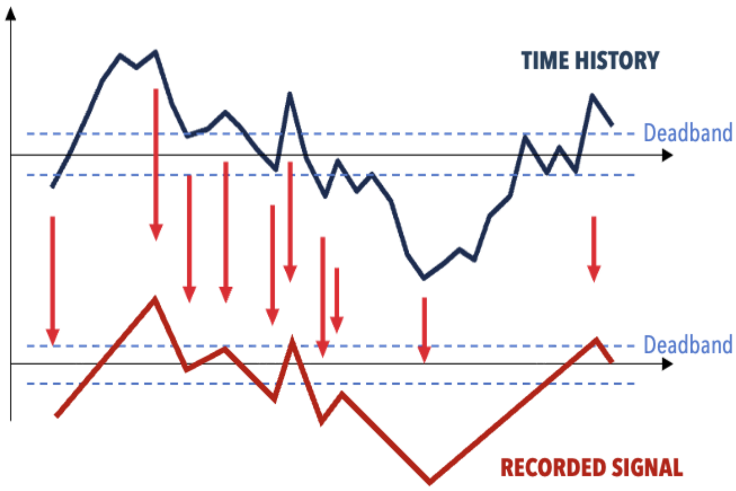

Fig. 4.5 PV recording principle applied to a sample signal

The AJT HUMS operates as a parametric system. This approach involves acquiring a limited set of raw data (around 50 signals), from which more complex and derived information is subsequently calculated. For example, rather than directly sensing aerodynamic forces through a specific sensor, these forces are computed based on parameters such as altitude, speed, attitude, etc. A total of 256 different parameters are saved in the LOAD file including a wide range of signals (both as sensed and calculated from the raw signals): LOAD factors, linear speeds, angular speeds, angular accelerations, attitude angles, Equivalent Airspeed (EAS), Calibrated Airspeed (CAS), pressure altitude, Mach number, LOADs, etc. One txt LOAD file for each flight is produced and saved, resulting in an enormous number of files to be analyzed.

For each flight, the HUMS generates a single LOAD text file and the data are systematically organized into folders. Specifically, a dedicated folder is created for each aircraft, containing all relevant text files, as illustrated in Figure 4.6. Each text file is labeled with the corresponding flight date, and within each file, the data are structured in a comma-delimited format. Finally, the initial lines of these files provide essential information about the flight.

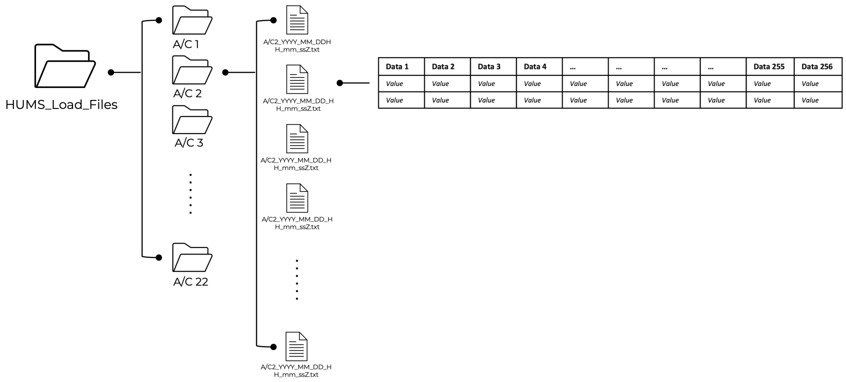

Fig. 4.6 Folder organization and hierarchy for HUMS LOAD files

HUMS systems are envisioned at a system level and they do not encompass the logging of equipment level data. Moreover, since HUMS was designed for SHM analyses, there was no need to record avionic bus signals or any other system level data. It is trivial that having this kind of data would have been extremely useful to understand and model equipment health. Nevertheless, Figure 4.7 shows the available signals related to the actuators: as can be seen from the graph, no actuator level data is saved in the HUMS LOAD file apart from the HT position, logged by the FCC but obtained synthetically from the jack position.

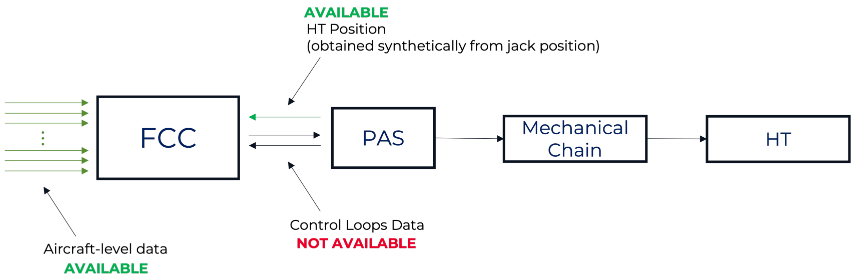

Fig. 4.7 Available signals in relation with the actuator.

4.1.2 PDF Documents

Although PDF documents such as technical documentation, TQ, and SI are primarily intended for consultation and knowledge acquisition and do not need to be treated as data, certain PDF documents require integration and management as data. "Equipment Registers" is an umbrella term in which the operational history of each equipment is reported. The creation of a coherent unique set of Equipment Registers (one for each actuator SN) required the assembly of information coming from different sources: cannibalization registers and loading/unloading registers. For instance, loading/unloading registers are in the form of hand-written PDF documents that have to be transcribed into machine readable csv files. Scheduled maintenance records are in the form of machine readable PDFs and have to be converted into csv files as well. URs have been assembled from transcribed Equipment Registers, as well as in-house UR tracking databases. OCR procedures have proven to be unreliable in these cases due to variability in individual handwriting, increasing the probability of errors. Consequently, manual transcription has been carried out.

The data acquisition activities conclude the first module of the framework's first layer. This activity, although it may appear to be a simple data collection phase, is crucial, as the quality of the data is fundamental to the success of the PHM system. Not only high-quality data ensures accurate analysis and decision-making, but also provides significant methodological value for similar projects. In fact, there are very few examples in the literature that highlight this critical phase, underscoring its importance in developing effective PHM systems [64].

4.2 Data Pre-Processing Module

Data pre-processing is critically important to ensure that data are meaningful and usable, effectively transforming potential insights into actionable information [110]. More often than not, good information is immersed in a noisy dataset that cannot be processed as is as it would lead to macroscopic errors and data artifacts. In other words, data pre-processing involves transforming and manipulating raw data to improve its quality, consistency, and guide future analysis.

This project has to deal with real life data collected during the asset operation. Data obtained from real-world systems are often different from those logged during experimental setups or simulations, often from many perspectives. In controlled environments, data acquisition frequencies can first be set and then synchronized to capture sensor logs simultaneously, and sensor configurations can be optimized to improve data quality. The amount of data collected is mainly influenced by computational resources, the duration of the experiment, and the sampling frequency. Currently, to our knowledge, there are no comprehensive and up-to-date protocols or guidelines for managing, organizing, and preparing HUMS sensor monitoring data intended for training and implementing PHM frameworks. Moreover, the steps involved in the pre-processing phase can vary depending on the available data and the chosen strategy, making data curation and preparation often a case-specific process rather than following a pre-made methodology. On the other hand, this should not surprise as data are often very platform specific and each case requires a specific protocol and specific techniques which are not transferrable. Many studies report their custom made pre-processing pipelines which can be extremely useful but, in the end, each case must be treated separately. In this case, there is no "one size fits all" and data have to be explored and analyzed to highlight specific inconsistencies or noisy values that could interfere and hinder the methodology results. Commonly performed actions can be grouped up with the terms: data cleaning [111], data cleansing [112], data transformation, feature pre-selection or selection [113], data normalization/standardization [114], etc.

In order to provide a sequential description of the activities performed in the Pre-Processing Module, the remaining of the section is divided according to some macro areas.

4.2.1 HUMS LOAD Files

Data Cleaning and Cleansing

Both maintenance and usage data may include inaccuracies, redundant information, errors, or noise. In fact, errors in usage data may arise from sensor malfunctions or input error by the ground crew before departing (e.g., setting the wrong fuel weight). Therefore, conducting a sort of check on some known column is crucial. The entries in the LOAD file are categorized as either "ground" or "flight" rows, reflecting whether the data was recorded while the aircraft's systems were operating on the tarmac or during flight. Consequently, ground rows have been excluded from further consideration. In addition, instances of NaN (Not-a-Number), missing value, and zero values have been excluded from the analysis to ensure data integrity and consistency. Some flights contained general corrupted information and were therefore excluded.

Data Type Conversion

Although data can generally be considered safe and reliable after the data cleaning process, it may still be necessary to adjust the data type of certain variables. Converting categorical data, such as switching from string to numeric formats or vise versa, is essential for effective data analysis, as many packages, algorithms, or functions handle only certain types of variables and may throw hard to decode exceptions otherwise. During the data pre-processing phase, significant attention was devoted to establish a coherent data structure and standardized formats. As already said, this approach is central to ensure that the data can be easily accessed and effectively analyzed in subsequent stages in a single framework.

Date and Time Formatting

The date and time for each flight, along with the timestamps for each recording, have been formatted as date-time objects and Unix timestamps, which is a widely used format for storing time values. Unix timestamps are an extensively used method for representing time in applications that require date and time manipulation. Unix timestamps are defined as the number of seconds that have passed since the Unix Epoch, which is set at January 1, 1970, at 00:00:00 UTC. This system provides an absolute, consistent, and unique way of tracking time across different platforms and programming languages, exploiting their ability to represent time in a format that is independent of time zones. The adoption of Unix timestamps in this project has been decided to eliminate any possibility of misunderstanding between similar data conventions (e.g. dd/mm/yyyy and mm/dd/yyyy) which could lead to catastrophic error in the next phases of the project. There are Python packages which enable seamless shift between time formats. The same approach has been adopted for URs and every other documents containing time and date information.

Feature Pre-Selection & Dimensionality Reduction

As discussed in Section 4.1.1, the total number of signals recorded by HUMS is 256, which represents a vast pool of signals to use in the analysis. Drawing on the research group's prior experience, expert input, and findings from the literature, it was concluded that extracting all 256 variables was not necessary for this study.

The rationale for this decision is that numerous variables may not be relevant to the case study and that including them could unnecessarily enlarge the dataset and slow down calculations. This approach is in line with the variable selection methodologies proposed by Ten Zeldam et al. [115] and Schoemakers [64].

Based on physical reasoning and flight mechanics knowledge, from the original 256 signals, a total of 50 variables have been pre-selected along with additional contextual information (e.g., Flight ID, UTC Date & Time, etc.). These 50 variables are identified as potential feature candidates that can be associated with the health status of the AJT HT PAS. Examples include aircraft variables within the longitudinal plane, the positions of mobile surfaces, and the forces acting on the tail and HT. This step has been presented in two conference papers [5, 6].

Synthetic Features

After a thorough analysis of the available signals and feedback from flight mechanics and actuation system experts, it was decided to combine four signals together to create two additional features that could represent and track actuator health in a more precise way. The two added signals are the mechanical works for each HT actuator (right and left) obtained using the information on the position of the actuator and the moment applied on the surfaces. The mechanical work can be representative of the overall usage rate of the actuator since it is linked to both the movement of the actuator jack and the forces it produces. Mechanical work for the right HT for the "i-th" row has been obtained as follows:

$$Work_{i,Right_{HT}} = M_x \cdot \frac{\Delta\alpha_{i,Right_{HT}}}{b} \tag{4.1}$$

where

$$\Delta\alpha_{i,Right_{HT}} = \alpha_i - \alpha_{i-1} \tag{4.2}$$

$\Delta\alpha_{i,Right_{HT}}$ refers to the HT position difference, calculated for each time step by subtracting the value in the previous row (*i-1*) and the value in the next row (*i*). b is a constant arm of 1 meter. As a result, considering the two signals linked to the position difference at each time step and the two mechanical works, the final number of signals is 54, complemented by complementary information. Table 4.1 reports the

Table 4.1 List of complementary information (a,b,c,d) and flight parameters (1-54) Flight dynamics, loads and aircraft level signals are logged.

Number	Flight Parameter (FP)	Notes
a	Mode	Flight/Ground
b	Flight ID	Increasing Number on each FCC power up
c	Counter	Increasing counter
d	UTC Date Time	DD_MM_YYYY – hh_mm_ss
1	Weight	[kg]
2,3,4,5,6,7	LHT Fx, Fy, Fz, Mx, My, Mz	Six load components acting on the left HT. [N or Nm]
8,9,10,11,12,13	RHT Fx, Fy, Fz, Mx, My, Mz	Six load components acting on the right HT. [N or Nm]
14,15,16,17,18,19	Tail Fx, Fy, Fz, Mx, My, Mz	Six load components acting on the last fuselage section. [N or Nm]
20,21,22	Nx, Ny, Nz	z, y, x axes accelerations [g]
23,24,25,26	Buffet Coefficients	-
27,28,29	p, q, r	Roll, pitch and yaw rate [deg/s]
30,31,32	p, q, r dot	Roll, pitch and yaw acceleration [rad/s^2]
33,34,35	V (NEU)	Speed North, East, Up [ft/s]
36,37	AoA, AoS	Attack and sideslip angles [deg]
38	LEF Deflection	Leading Edge Flap Position [deg]
39,4	Stick Position (long, lat)	Longitudinal and lateral stick position [mm]
41,42	LHT, RHT deflection	Left and right HT position [deg]
43	Pedal position	[mm]
44	Flap position	Trailing Edge Flap position [deg]
45	Pressure altitude	[ft]
46	TAS	True Air Speed [Knots]
47	Mach	Mach number
48,49	Left and right throttle	[%]
50	Static air temperature	[°C]
51,52	LHT, RHT deflection variations	Left and right HT position difference with the previous value [deg]
53,54	LHT, RHT Work	Mechanical work of the actuators (calculated)

list of Flight Parameters (FP) and complementary information indicated as (a,b,c,d) along with the unit of measurement and some additional notes if needed.

General Data Operations

The term General Data Operations, in the context of this thesis, is used to refer to data structure and management operations to create a more optimized data structure. Under this umbrella, it is possible to identify the following:

- Flight Information Processing: The flight number and date for each LOAD file are identified from the initial rows of each HUMS file and recorded accordingly.
- Flight Number Grouping: An additional column has been introduced to include a sequential flight number, facilitating easy retrieval of each file in the database.
- Concatenation of files: Files pertaining to each aircraft have been combined into a single data frame.

- Column Remapping: Columns have been reorganized and renamed.
- Saving: data frames have been saved in dictionaries as well as locally in .csv and .json files.

4.2.2 Equipment Registers

In the context of this thesis, the term "equipment registers" identifies the logistic documentation which reports the history of the fleet of actuators, including on which aircraft TN they were mounted, for how many flights, why they were unloaded as well as the operational history of the component (i.e. when the equipment has been sent to the OEM and the return date). Equipment registers are equivalent to the documentation that is commonly called a "Log Card" and that follows the equipment along its operational life. Equipment registers are often based on a paper document that travels along the PAS to the OEM, to the operative base, and to the warehouse. As a result, the pre-processing step for Equipment registers involved the hand transcriptions from paper tables to a csv-based file of the records. Since the documents were signed and compiled by different entities, technicians, and with different styles, automated procedures were found to be inefficient, leaving hand transcription the only viable option. As a result, a .csv document has been created with the following information:

- Equipment SN
- Aircraft TN (loading and unloading)
- Reason for change (if present)
- OH of equipment and aircraft at the unloading time
- Date of change

Much effort has been put into gathering information from different sources and merging them into a single database. A further step was required to identify the last flight before the change, as well as the conversion of date and time indications into Unix timestamps. All of this information will be referred to as Configuration Management (CM) throughout the remainder of the thesis. Equipment Registers should in principle represent the primary source of information on possible fault modes; however, as further discussed in the next Section, their content is not sufficiently consistent or statistically relevant to support the development of a reliable fault-detection strategy.

Fault Mode Statistics

As already stated in Section 1.10 and 2.1, this framework targets the prediction of Unscheduled Removals (URs) of the AJT HT EHA, rather than the identification of specific actuator fault modes. The focus is therefore on anticipating no-go conditions that affect asset availability and fleet management, while classification of the exact removal cause is outside the scope of the approach.

As already mentioned in Section 4.2.2, the unloading reasons are sometimes mentioned in the Equipment. However, they are often not coded and written by hand, leading to a not standardized unloading code formulation. Following an analysis with SMEs, the initially broad range of documented causes was refined to a limited number of primary unloading reasons, consistent with existing literature on EHA fault trends.

Literature on EHA indicates that actuator seal wear frequently leads to internal leakage between chambers. Seal degradation is commonly modeled as a cumulative damage process that progresses over time, representing a gradually developing fault. This approach supports the concept that certain faults evolve slowly into failures, and their degradation can be tracked within the PHM framework by utilizing cumulative metrics that incorporate flight history and duration.

However, due to limited data availability—specifically, the absence of hydraulic pressure measurements and other relevant data in the AJT HUMS—modeling the seals' degradation on the actuator piston is currently not feasible.

Figure 4.8 presents a histogram of fault modes, illustrating various failure types; specific labels are omitted for confidentiality reasons.

Detailed analyses on the possibility of treating each failure code separately in order to improve prediction accuracy were carried out as ablation studies and are reported in Section 5.4.1, showing that the limited statistical base does not support any tangible benefit from modeling different fault modes independently. As the fleet accumulates additional flight hours and clearer trends emerge, such a distinction may

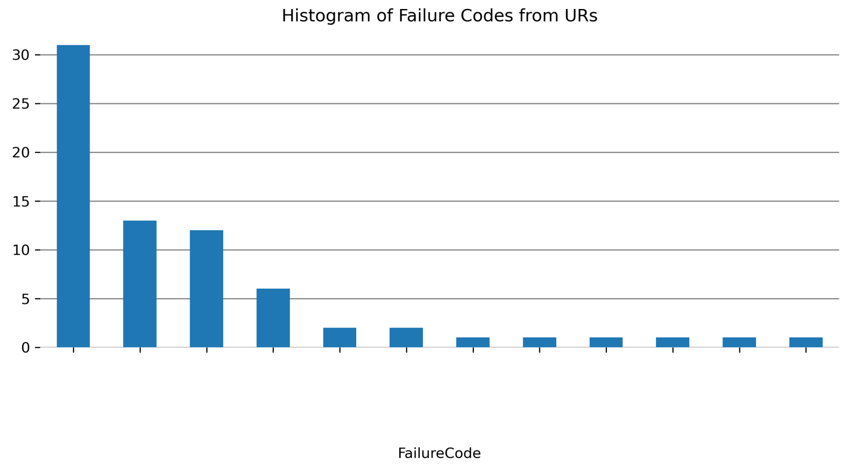

Fig. 4.8 Histogram of Failure Codes from URs with omitted labels.

become more meaningful, potentially improving prediction accuracy and opening the possibility of classifying URs according to their underlying cause.

4.2.3 URs

The pre-processing phase for URs has been quite similar to the one related to Equipment registers. UR-related information was found to be scattered in various documents and reports, and much effort has been put to create a coherent and comprehensive database for the entire aircraft fleet. As a result, a single ".csv" file has been created that gathers the fault code (if present), the date of the UR, and the OH of the equipment. As per the equipment registers, further steps were required to identify the last flight before the UR as well as the conversion of date and time indications into Unix timestamps.

4.2.4 General documentation

General and additional documentation was retrieved in a wide range of formats (e.g., ".pdf" or excel documents, ".xls" or ".csv" files, etc.). The information retrieved has been used in the framework development when needed and to inform the framework choice.

Chapter 5

Layer 2: Data Processing, Features, and Prognosis

5.1 Feature Engineering Module

The Feature Engineering Module is one of the four modules built within the Data Processing, Features & Prognosis Layer. The module inputs are the FPs, the CM files, FHs and URs. The Feature Engineering Module is responsible for the creation and selection of a set of Cumulative Features (CFs) which can represent the evolution of the equipment health in time and provides the parameters of the hazard Probability Distrbution Functions (PDFs) which will be then exploited in the Probabilistic Health Predictions Module.

Following the blocks reported in the main scheme of Figure 3.2, the first action performed by this module is the Statistical Moment (SM) calculation. SMs are a common way to lump the characteristics of a signal or of a data distribution in general. In other words, SMs are quantitative metrics that describe the characteristics of a certain signal or distribution. They provide insights into the shape, spread, and central tendency of data sets. Among the most famous SMs, the first four are the most popular as they are able to merge the statistical contribution of the signals into a handful of parameters.

If a random variable X is considered, the n-th moment about the origin is given by Eq. 5.1:

$$\mu_n' = E[X^n] = \int_{-\infty}^{\infty} x^n f(x)\,dx \tag{5.1}$$

where $f(x)$ represents the X PDF. A subcategory of SM is represented by central moments, which are calculated relative to the mean of the data distribution.

$$\mu_n = E[(X-\mu)^n] \tag{5.2}$$

where μ is the mean of X. Thanks to the generalized formula provided in Eq. 5.2, the first four central moments can then be formulized:

1. First central moment: Mean. The mean value represents the central tendency of the data and underscores that the aircraft or actuator has experienced high or low loads focusing on the baseline rather than the extreme values and tails [116]. Despite being a trivial statistical metric, it can provide useful insights on the forces acting on the actuator during the flights.

$$\text{Mean}(\mu) = E[X] = \frac{1}{n}\sum_{i=1}^{n} x_i \tag{5.3}$$

2. Second central moment: Variance. Variance indicates the dispersion or spread of data points around the mean. A high variance observed within a specific feature time series suggests that the aircraft has encountered values that differ significantly from the average. For instance, elevated variance in work values may indicate that the PAS jack has undergone a broad spectrum of movements and positions. Even if variance is a very common measure in statistics, it can provide useful insights on the spectrum of regimes encountered and stress. For instance, the authors of [117] used variance analyses on cumulative signals in their study on inverter fault detection.

$$\text{Variance}(\sigma^2) = E[(X-\mu)^2] \tag{5.4}$$

3. Third central moment: Skewness. Skewness quantifies the probability distribution asymmetry. A longer right tail leads to positive skewness, whereas negative skewness highlights a longer left tail. Variance has been extensively used in PHM and diagnostic applications for rotating mechanical equipment

[118]. High skewness can indicate values that consistently deviate from the mean within the load spectrum, highlighting their importance.

$$\text{Skewness} = E\left[\left(\frac{X-\mu}{\sigma}\right)^3\right] \tag{5.5}$$

4. Fourth central moment: Kurtosis. The fourth central moment is one of the most widely used basic statistical indicators for the diagnosis and prognosis of rotating equipment spanning from shaft to wind turbines due to the ability of kurtosis to concentrate and highlight anomalies in distributions [119, 118]. In particular, kurtosis shows the "tailedness" or "peakedness" of the distribution.

$$\text{Kurtosis} = E\left[\left(\frac{X-\mu}{\sigma}\right)^4\right] - 3 \tag{5.6}$$

   High kurtosis suggests heavy tails and a sharp peak, while low kurtosis indicates light tails and a flatter peak. A relevant literature application involving kurtosis is proposed by the authors in [120] where kurtosis has been applied for EHAs PHM, proving the approach robustness.

In conclusion, SMs provide essential insight into data distributions by quantifying central tendency (mean), variability (variance), asymmetry (skewness), and peakedness (kurtosis). The choice of relying on SMs is not arbitrary: it is driven by their high degree of explainability, their compatibility with the available data granularity, and their suitability for modular integration within the proposed PHM framework. Each moment can be directly interpreted by domain experts in terms of physically meaningful concepts such as average load level, variability of operating regimes, prevalence of extreme events, and concentration of stress cycles, which facilitates both validation and adoption in an industrial context. Moreover, given that the available information is at the aircraft level and is often affected by heterogeneous operating conditions and measurement noise, more complex high-dimensional descriptors (e.g., time–frequency or deep-learned features) would be harder to identify robustly and would significantly increase the risk of overfitting without a proportional gain in interpretability. Statistical moments can be defined for any order (even higher than four), but the first four are adopted here as a well-established set that captures the main distributional characteristics relevant for degradation analysis, while keeping the feature space compact enough to be propagated through the subsequent Monte Carlo and hazard models. This choice preserves the modularity of the framework, since additional families of features (higher-order moments, quantiles, frequency-domain indicators, etc.) can be incorporated in future extensions without altering the core architecture of the PHM pipeline.

The four SMs have been applied on the pre-processed data set for each flight. As a result, each FP of each flight is described by a total of four moments. Therefore, a flight is described by a total of 216 values: 4 SMs applied on each of the 54 FPs. This process is represented in the scheme in Figure 5.1. Each aircraft of the fleet has a set of FP files, one for each flight containing the time series of the 54 pre-selected FPs. For each flight, the SMs are calculated for each FP, resulting in 216 values for each flight. This process is repeated for each flight of each aircraft, providing a way to track the usage of the system from flight data.

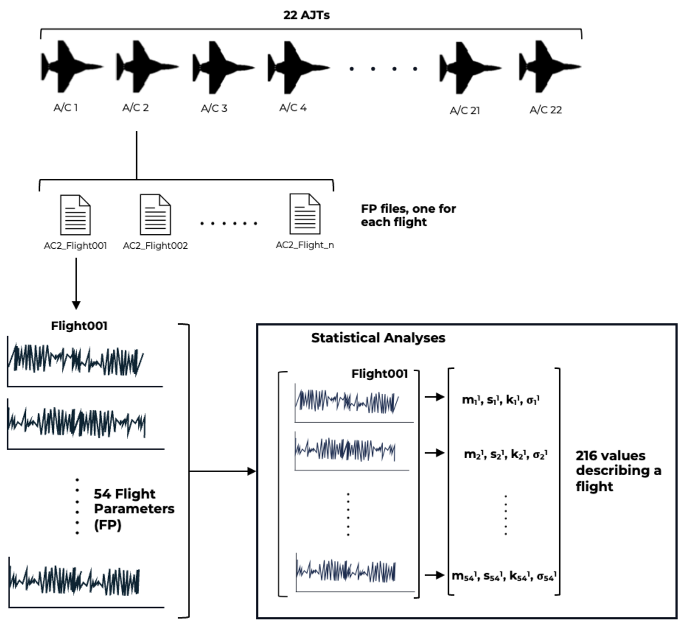

Fig. 5.1 SM creation process from FPs.

Following the scheme reported in Figure 5.2, the next step involves the assembly of PAS data. This activity is central in following the PAS operational history. Aircraft components are designed by OEMs for high maintainability, using the concept of Line Replaceable Units (LRUs) to facilitate maintenance operations and troubleshooting. The concept of LRU describes a component that can be easily and quickly replaced during maintenance operations without the need for extensive disassembly of the surrounding system. This feature is crucial to minimize downtime and ensure operational readiness in complex operational environments. As a result, equipment can be exchanged during service for reasons such as cannibalization, troubleshooting, or operational readiness requirements. Moreover, when an aircraft component fails and is unloaded from an aircraft Tail Number (TN) and sent to the OEM for repair, there is absolutely no assurance that the repaired component will be reinstalled on the same aircraft TN from which it was removed. Instead, it may be added to the general spare parts inventory and/or installed on a different aircraft TN upon its return from the OEM according to operational needs. This concept, which definitely improves fleet maintainability and readiness, on the other hand, makes health tracking much more challenging.

FPs are tracked and saved by aircraft; therefore, it is necessary to gather the operational life of PAS by "cutting and pasting" relevant FP trends from the TN of the aircrafts on which they were loaded through the CM files. In fact, tracking the status of a specific piece of equipment, identified by its Serial Number (SN), is not equivalent to monitoring the operational history of the aircraft, identified by its TN, on which the equipment was initially installed. On the other hand, tracking what the equipment has experienced is essential to understand the way that a particular flight or mission affected the actuator health. Although it is more a management issue, this problem is complex and requires a comprehensive study and analysis of CM information. However, without establishing a link between SN and TN, it becomes impossible to utilize aircraft operational data to predict the future performance of the PAS. Figure 5.2 shows an exemplified assembly history of a specific serial number, PAS 002. According to the exemplified equipment register, this unit was installed successively on aircraft 1, aircraft 2, and aircraft 3. This information made it possible to reconstruct the operational life of PAS 002 by extracting the data generated by each aircraft during the periods in which the unit was installed. These data segments were then combined in chronological order to reconstruct the complete operational history of that specific PAS.

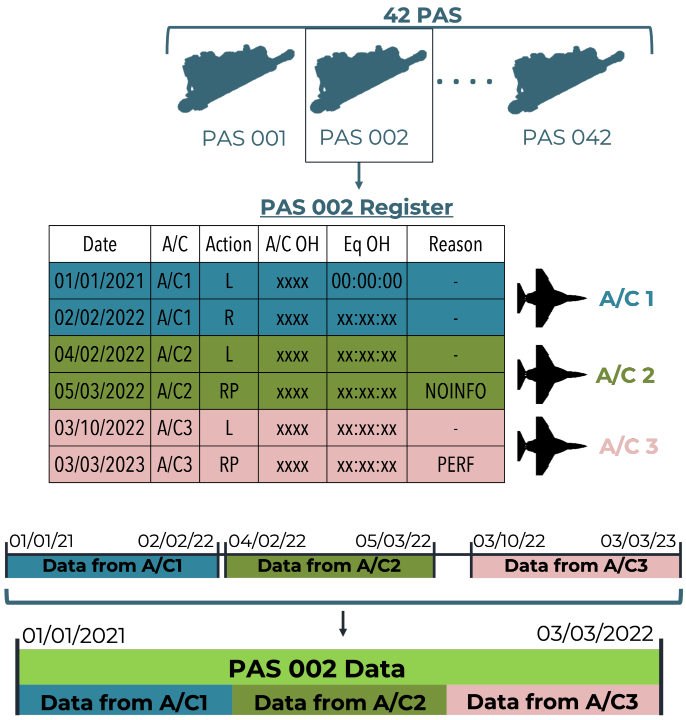

Fig. 5.2 PAS data assembly process. The PAS 002 data is assembled from aircraft 1, 2, and 3, organized by loading (L) and unloading (R - Replacement or RP - Repair) dates. Occasionally, the reasons for equipment unloading, such as "PERF" indicating a performance test failure, are documented.

At this stage, since the SM data for each flight are available and organized for each PAS, the Condition Indices (CIs) can be assembled by merging the knowledge of what happened during the flight (SM data) with the knowledge of how long the flight was (FHs data). The concept behind the creation of cumulative trends has already been used in the PHM community as a way to track the evolution of equipment health over time in a completely explainable way. In the case of this project the choice of using this approach is strictly linked to the granularity of data at hand and it is deemed the best way possible to integrate aircraft level data to predict PAS level health. Other applications of cumulative trends are proposed in [117] where the fourth statistical moment is considered as well as in [121], where the authors applied a cumulative of the fourth moment to identify working regimes. The result of this step is a set of 216 cumulative CI, which describe the life of the PAS both in terms of what happened during each flight (thanks to the information contained in the SM) and the duration of each flight (thanks to the duration FD).

The process of creating CIs is explained in Eq. 5.7 and Eq. 5.8.

$$CI_{i,j}(f) = \sum_{k=1}^{f} SM_{i,j}(k) \cdot FD(k) \tag{5.7}$$

$CI_{i,j}(f)$ indicates the CI at flight f related to the SM $SM_{i,j}$. $SM_{i,j}(k)$ indicates the SM i applied on FP j with respect to flight k. i is an index which may assume "mean", "variance", "skewness", or "kurtosis". j ranges from 1 to 54, indicating the various FPs. For example, $CI_{mean,1}(k)$ is the way the CI value obtained using the SM "mean" on the FP number 1 related to the flight k is mathematically formalized. On the other hand, $FD(k)$ represents the duration of the flight k.

A vectorized formula can also be defined, referring to the vector $\overrightarrow{CI_{i,j}}$, as shown in Equation 5.8:

$$\overrightarrow{CI_{i,j}} = cumsum(\overrightarrow{SM_{i,j}} \cdot \overrightarrow{FD}) \tag{5.8}$$

where $\overrightarrow{CI_{i,j}}$ indicates the CI of $\overrightarrow{SM_{i,j}}$. $\overrightarrow{SM_{i,j}}$ is the vector that contains all concatenated FP j calculated thanks to the SM i, and $\overrightarrow{FD}$ is the vector of flight durations.

In other words, for each flight, the SM value is multiplied by the flight duration (FD), and the resulting increase is subsequently added to the previous cumulative value. In essence, CIs are derived by integrating the SM over time in a finite integral fashion; specifically, this involves multiplying the flight-related SMs by the FD to create a trend that accurately reflects the aircraft's usage. Thus, each flight is defined by a sequence of CI increments (one for each FP). This approach has been systematically applied to each PAS and each SM.

The only data source not yet exploited is the record of URs. It should be recalled that the objective of the framework is not the identification of specific fault modes, but the prediction of UR themselves, as events affecting actuator availability and fleet management. Once the CI trends have been calculated, their evolution is analyzed in conjunction with the URs recorded in the UR register. By examining the CI trends alongside the URs, we can assess the magnitude difference between consecutive URs or between the initial flight and the first UR. Following these analyses, histogram representations were utilized to visually depict the distribution of the derived intervals. These intervals represent the variations in the value of a specific CI over time, providing valuable information for identifying the most significant CIs.

Figure 5.3 illustrates the process for a specific actuator, identified as PAS SN 002. The graph on the right displays a particular CI and employs a color code corresponding to the aircraft on which the PAS was installed, based on the date. The table in the upper right corner of the figure shows that this specific PAS has flown for a total of 1145 flights. The CI is highlighted by several red vertical dotted lines, which indicate the removals that PAS 002 has undergone. Variations are calculated between the beginning of the PAS operational life and the first UR), as well as between subsequent URs. Notably, PAS 002 experienced three URs; however, only two of these were due to failures, marked as "RP" for Repair. The initial UR was classified as a Removal (R), which can be linked to cannibalization or to the need on unloading the PAS to load it on another aircraft TN.

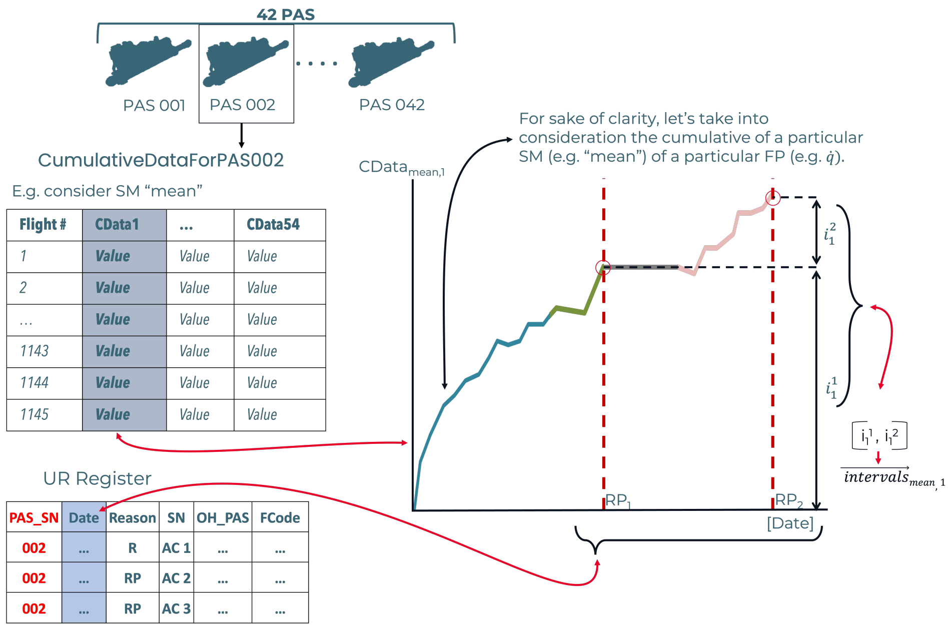

Fig. 5.3 An example of a calculation of intervals values for PAS 002. The plot is a exemplified sketch of a possible trend for explanatory purposes.

The histograms obtained have been used in two different ways:

1. **Identifying Periodicity**: Histograms that display the intervals between URs are a way to identify CIs that show periodic behavior and a relatively stable value between two URs. Ideally, the optimal CI would show consistent changes in magnitude between URs, represented by a histogram with a single bin. Each histogram is accompanied by a normal PDF curve, defined by its mean

and standard deviation. This setup enables the calculation of a normalized Signal-to-Noise Ratio (S/N Ratio), which ranks the features according to their informativeness. Specifically, the most desirable features are those with the highest S/N ratios and the narrowest histograms. A total of five top CIs were determined to be sufficient to characterize the health of the PAS: these include the skewness applied to the mechanical work on the right actuator, the skewness of the north speed component V_n, the mean of the last section of the fuselage moment M_z, the moment M_x applied to the right horizontal tail calculated using the mean and the kurtosis of the east speed component V_e. The most representative CIs are called Cumulative Features (CFs)

2. **Hazard Distribution Representation**: Histograms also serve a central role in the subsequent analysis steps. They not only help identify the most significant features, but also represent the hazard threshold distribution. If a histogram indicates a change in magnitude between two URs, this suggests that reaching this change increases the probability of an UR occurrence. Therefore, when fitted with a custom PDF, histograms effectively illustrate the distribution of hazard thresholds. In fact, as stated by Acuña et al. [122], one way to define a threshold function is using post-mortem statistical data obtained from a fleet of running equipment. Thanks to the values of these condition indicators at the recorded failure instances, a parametric or non-parametric hazard function can be developed by considering the cumulative function of the obtained PDF.

The histograms of the 5 best CFs are shown in Figure 5.4 together with the normal PDF to rank them. It is important to clarify that the normal distribution fits shown in Figure 5.4 are used solely for the calculation of the S/N Ratio. The normalized Signal to Noise ratio (SN Ratio) has been calculated as shown in the Eq. 5.9:

$$S/N_{norm} = \frac{norm(m)}{\Sigma} \tag{5.9}$$

where m and Σ are the mean and standard deviation of the fitted normal distribution.

This metric requires the extraction of mean and variance and therefore the same distribution model must be applied uniformly to all features to ensure a fair comparison. Although this step provides a consistent basis for feature ranking, it is not intended to reflect the true underlying distribution of each feature. In the

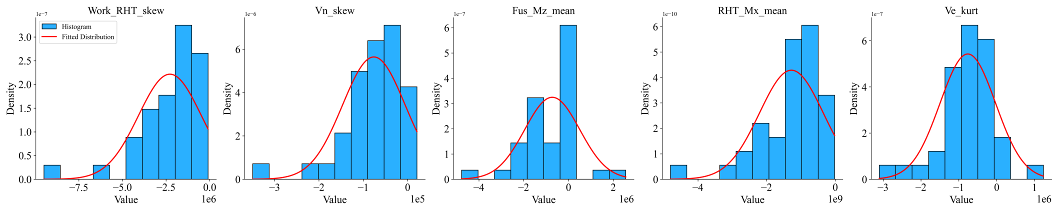

Fig. 5.4 Histograms of the best 5 CFs.

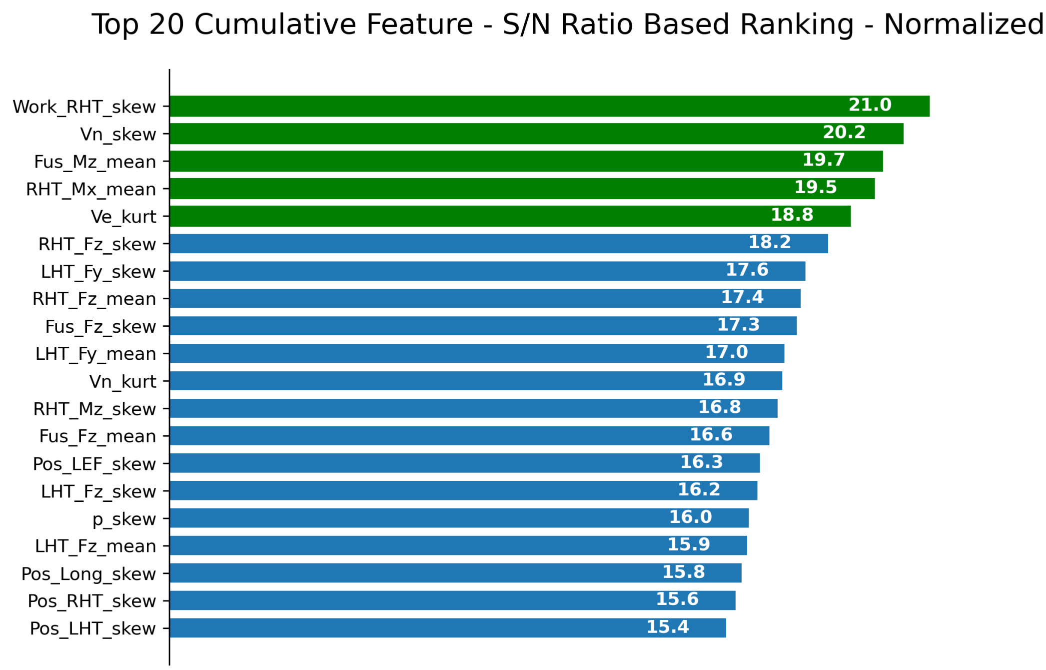

Fig. 5.5 Ranking results based on normalized S/N ratio. The top 5 CFs are colored in green.

subsequent sampling stage, the CF distributions are instead fitted using a more refined and feature-specific approach.

The ranking results are reported in Figure 5.4 along with the normalized S/N ratio value in Figure 5.5. It should be noted that the most informative CFs are physically connected to the mechanical behaviors of the HT and their corresponding impact on the AJT flight mechanics. Additionally, the observed prevalence of skewness, mean and kurtosis, coupled with the lack of variance, may indicate a more significant influence of higher loads in relation to mean value.

The two module outputs are hence the Top 5 CFs trends obtained from the S/N ratio ranking and the hazard probability distributions for each CF. Each hazard distribution has been analyzed and fitted with the optimal PDF utilizing the "Distfit"

Python library, which facilitates probability density fitting of univariate distributions through both parametric and non-parametric methods.

In the parametric approach, the Distfit library evaluates a total of 89 theoretical distributions to identify the best fit. The goodness of fit for these PDFs is assessed using the Residual Sum of Square errors (RSS), allowing for the selection of the most representative distribution. The considered potential parametric PDF function ranges from common formulations (e.g., Weibull, Gamma, Cauchy) to advanced formulations (e.g., Log-gamma, Johnson SU, t-Student). In other words, this approach allows for the selection of a specific parametric PDF function, ensuring a high level of customization and accuracy. The RSS quantifies the divergence between the predicted and observed empirical data values, reflecting the discrepancies in the estimates. It serves as a measure of the difference between the data and a given estimation model. The RSS is calculated using the formula reported in Eq. 5.10:

$$\mathrm{RSS} = \sum_{i=1}^{n} (y_i - \hat{y}_i)^2 \tag{5.10}$$

Some further outlier removals are performed on the distributions to remove outliers that could hinder the definition of a clear hazard threshold distribution. The values provided represent criteria for removing outliers from the various distributions:

- The first criterion involves "Work_RHT_skew" and indicates that any values less than or equal to -6 million are considered outliers and are removed from the dataset.
- The second criterion relates to "Vn_skew" and specifies that any values less than or equal to -250,000 are removed.
- The third criterion indicates that any bin of the CF "RHT_Mx_mean" must be bigger than or equal to -4 billion.
- The fourth criterion on "Ve_kurt" means that any values greater than or equal to 1 are removed from consideration.
- The fifth criterion specifies that any values less than or equal to -50 million are excluded from the dataset of the CF "RHT_Fz_skew".

In general, these filtering criteria aim to enhance the quality of the dataset by removing extreme values that could skew the results or lead to incorrect conclusions in future analyses. By setting these thresholds, only data points within reasonable limits are retained for modeling and evaluation, ensuring a more reliable interpretation of threshold metrics.

5.1.1 Additional Analyses

In the context of the Feature Engineering Module, a wide range of parallel analyses and studies have been carried out to support and inform the next steps. In the following subsections, the main results are reported and explained.

Sampling Analyses

A sampling analysis has been conducted for four aircraft. Sampling analyses are important to identify the specific signal bandwidth and to understand if the sampling is high enough to capture flight or actuator dynamics. The results are consistent across all aircrafts analyzed and can therefore be generalized to the entire fleet, as recording equipment, training flights, and overall subsystem infrastructure are utilized uniformly throughout the system. An example of the analysis for an individual aircraft is presented in Figure 5.6. The sampling analysis reveals a highly variable frequency, with only a small percentage (up to 22%) of irregular and sparse batches of high-frequency data, while the majority of samples are collected at frequencies below 5 Hz. It is important to remember that, as already explained in Section 1.10 and 4.1.1, the HUMS mounted on the AJT was originally designed for SHM applications rather than PHM, and as such, high-frequency sampling was not a requirement. Moreover, as already explained in Section 4.1.1, a peak-valley recording principle had been adopted to log data in the AJT HUMS and, as a result, the data are saved only when a peak or a valley is exceeded for one of the monitored parameters. In conclusion, while during test flights additional sensors and logging equipment are mounted on the platform, no constant/high frequency sampling logging is carried out on production aircraft. The results of the sampling studies highlighted that, with the current data sampling strategy, it is impossible to track the dynamics of the aircraft and actuator as well. In fact, common aircraft dynamics bandwidths are in the range of 3-4 Hz for an F-16 Falcon, while actuator dynamics is often one order

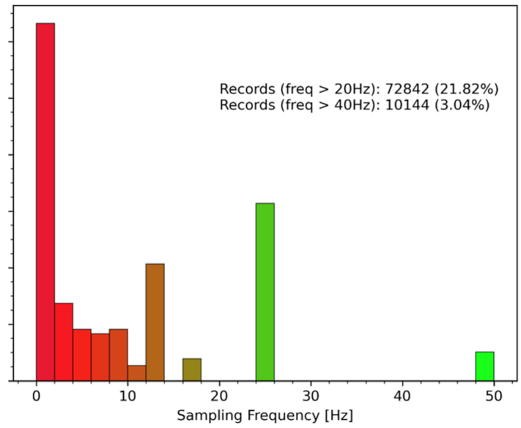

Fig. 5.6 Sampling analysis for one TN: most of the sampling is carried out at very low frequencies, therefore not allowing equipment level degradation studies.

of magnitude higher [123]. The Nyquist-Shannon sampling theorem states that if a system uniformly samples an analog signal at a rate that exceeds the signal's highest frequency by at least a factor of two, the original analog signal can be perfectly recovered from the discrete values produced by sampling. There are two important points in this definition: the first one is the mere "factor of 2" which is definitely not satisfied in our timeseries. The second important point is that the samples must be uniformly distributed. This is definitely not satisfied, as the PV recording principle leads to extensive instant where no data are available for extended minutes. As a result, raw time series cannot faithfully represent aircraft, dynamics, maneuvers, or actuator dynamics.

Recent trends in signal processing are increasingly focused on upsampling techniques applied to time series data to enhance data mining capabilities [124]. These techniques, often utilizing deep learning tools, typically train models on small batches of high-frequency data that can then be applied to low-frequency datasets. However, in this particular instance, we were unable to identify any reliable high-frequency batches. Furthermore, incorporating high-frequency data into an HUMS sequence without data may lead to spurious and unrealistic maneuvers. Given that empty sequences can last from seconds to minutes, there is no information on the trajectory of the aircraft during these intervals. Consequently, the application of upsampling techniques to improve data interpretation is deemed to be inapplicable, as the resulting data cannot be trusted to accurately represent aircraft or actuator dynamics.

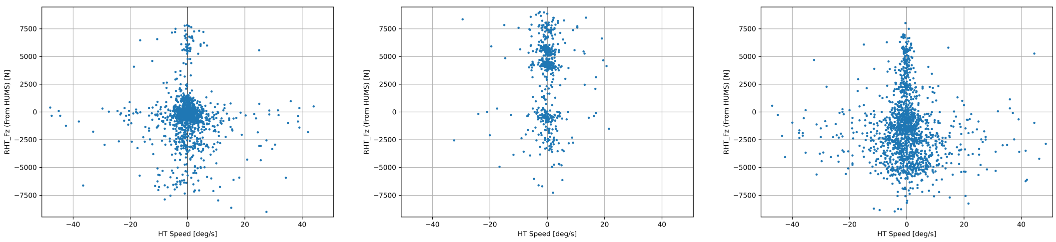

(a) Flight 1 (b) Flight 74 (c) Flight 138

Fig. 5.7 Speed–force working-point maps for three representative flights (Flights 1, 74, and 138) of the same PAS SN. No consistent evolution or separable trend is observed across flights, suggesting that a low-level physics fault-to-failure degradation path cannot be reliably identified from the available signals alone and motivating the adoption of alternative techniques.

Speed-Force Diagrams

Given the lack of actuator signals, as already mentioned in Section 4.1.1, a possible way to identify some degradation process in the actuator is the use of bidimensional maps that report for a subset of flights the working points of the actuator speed with respect to the force provided by the actuator at each time point. In particular, the speed is obtained by the derivative of the position trend, and the obtained signal is then filtered to exclude spurious data created by the very low sampling rate. The force on the HT is already saved in the FPs and has been used as is. This is carried out for each point in the time series data and the final graph should provide some indication on the presence of a fault-to-failure process. However, the plots do not appear to show any consistent trend because of the very low sampling rate and the lack of actuator data to follow. This analysis led to the conclusion that the identification of a degradation path through the available signals was not feasible and that other techniques were needed.

For instance, Figure 5.7 illustrates three representative speed–force working-point maps (Flights 1, 74, and 138) built from the available HUMS-derived information for the same SN PAS. Actuator speed is approximated as the time-derivative of the position trace and filtered to mitigate spurious effects induced by the low sampling rate, while the corresponding HT force is taken directly from the computed fault precursors. Each scatter point represents an instantaneous operating condition during the flight. As shown, the maps do not exhibit a clear or repeatable evolution across flights, confirming that the identification of a meaningful degradation trajectory from the available signals is not feasible in this dataset and motivating the adoption of alternative approaches.

Bath-tub curve

Bathtub curves are trends used in reliability engineering to represent the failure rate of a component, system, or asset over its lifecycle. Its name is inspired by the typical shape that it assumes for mechanical systems, as it highlights three distinct phases of failure probability. The first phase, often called the infant Mortality (Early Failures) phase, shows a decreasing failure rate. Literature indicates the cause could be traced back to manufacturing defects, design flaws, or improper installation. These issues decrease over time, and, as a result, the failure rate decreases as well, reaching a stable plateau of constant failure rate. This is the second phase, which indeed highlights random stress events or external factors which may lead to unexpected breakdowns during stable operational periods. Finally, the wear-Out (End-of-Life Failures) phase is characterized by a sharp increase of the failure rate as the components begin to degrade due to fatigue, corrosion, or cumulative wear. This phase underscores the age of the equipment that is approaching end-of-life. The described trend is usually associated with mechanical equipment, and the failure rate is often imposed as a design requirement and verified through chain relationship between components with data obtained through spreadsheets, reference manuals, parallels with similar equipment or experimental tests. On the other hand, purely electronic equipment or software reliability is often modeled using other models since this kind of equipment does not suffer mechanical wear.

Insights have been derived from the analysis of URs logistical data across the entire PAS fleet. Figure 5.8 illustrates the frequency of PAS URs in relation to the Operating Hours (OH) written on the register for each UR. The use of relative occurrence was essential to normalize the number of URs at specific OH values based on the total number of PAS that reached those thresholds. Attempting to apply this normalization would have resulted in a biased analysis because of the limited number of PAS serial numbers that achieve elevated OH values.

The graph indicates a pronounced downward trend that corresponds to the initial decline characteristic of a bathtub curve. This pattern suggests that the majority of recorded URs occur early in the actuator's life-cycle. As already stated, during the early phase of this curve, often referred to as the "burn-in period," the failure

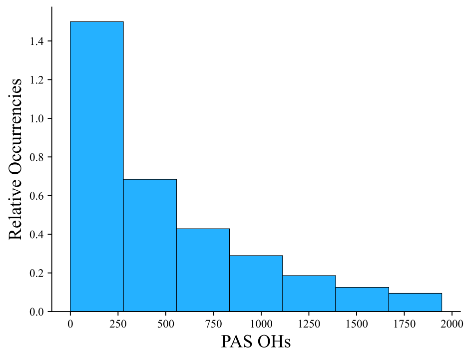

Fig. 5.8 Occurrences of UR events in function of the OHs reached at the UR event.

rate tends to be variable; it typically decreases due to various factors associated with phenomena identified as infant mortality in equipment registers. Consequently, these initial failures do not represent degradation, but rather stem from unpredictable malfunctions. Such failures are difficult to predict because of the lack of a fundamental degradation mode. This analysis highlights that the framework output may be more reliable as the fleet accumulates FHs, and a more solid statistical base will be achieved.

Finally, it should be noted that most of the excluded URs were located in the first bin, highlighting that those URs were probably caused by impromptu failures linked to infant mortality, loading issues, etc. These kinds of failure cannot be predicted due to the absence of an underlying degradation mode and hence have been excluded from the analysis.

Figure 5.9 reports the distribution of the number of URs experienced by each PAS serial number (SN) over its service life. Overall, the UR count per unit is low: no PAS exhibits more than three URs (and only one reaches three), while the majority of units experienced two, one, or zero UR events. This scarcity of failure/maintenance events further increases the difficulty of drawing statistically robust conclusions at the individual-SN level and reinforces the limitations already suggested by the bathtub-curve representation.

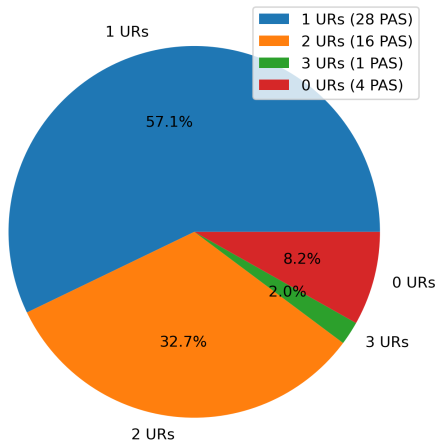

Fig. 5.9 Distribution of unscheduled removals (URs) per PAS serial number (SN). Pie chart showing the share of PAS units that experienced 0, 1, 2, or 3 UR events over their service life.

Flight statistics

An additional statistical insight was obtained from analyzing flight durations. Figure 5.10 presents a box plot that illustrates the flight times across the entire fleet. The analysis indicates an average flight duration of 58 minutes, consistent with training purposes of the AJT under study. The outliers observed in the data are associated with extended training sessions that involve external fuel tanks or transfer flights. In contrast, shorter flights are mainly attributed to test flights or brief training missions.

Data availability

To support the technical analyses aimed at identifying and deriving CI, a dedicated assessment of data availability was performed. Specifically, we generated data-availability timelines to visualize, for each PAS SN, the time intervals in which the different data sources are present or missing, and to highlight potential inconsistencies across databases like the one reported in 5.11.

The first timeline refers to the FH register: periods with no records are shown in red, the interval between the first and last available records is highlighted in light green, and the specific dates on which FH were recorded are shown in dark green. The second timeline represents the availability of HUMS files, using the same

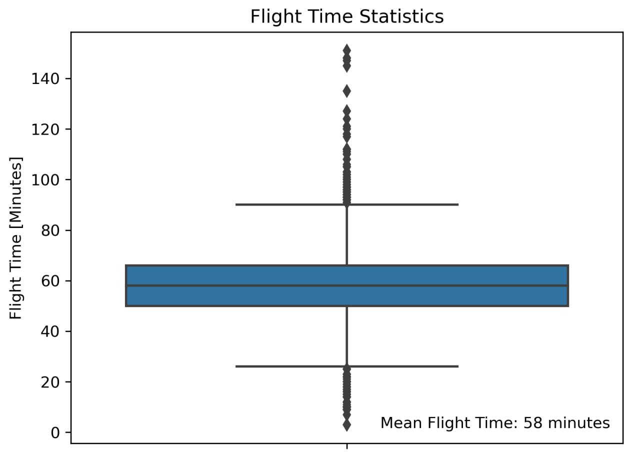

Fig. 5.10 Flight duration statistics.

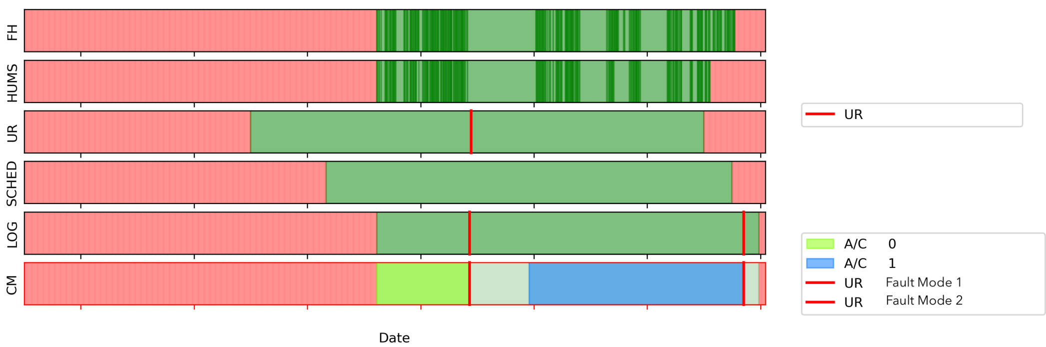

Fig. 5.11 An anonymized example of an availability chart.

convention (light green for the overall coverage window and dark green for recording dates).

The subsequent timelines summarize maintenance-related information. URs from operational databases are displayed as vertical event markers within the corresponding temporal span of the records. The availability of scheduled maintenance data is also reported to contextualize UR events within the planned maintenance activity.

Finally, two timelines are derived from the Equipment Registers dataset. The first depicts the Equipment Registers data coverage window (defined by the first loading date and the last update date) and overlays UR events, distinguishing different removal causes (e.g., failure-driven removals versus swaps) and reporting the associated failure code when available. The second LC timeline complements the previous one by reporting configuration-management information, i.e., the aircraft on which each PAS was installed and the corresponding installation periods, thereby enabling a clear reconstruction of the component operational history.

This type of analysis was systematically carried out for all PAS SNs, with the dual objective of (i) identifying data gaps and cross-source inconsistencies that may bias FP extraction and subsequent PHM modeling, and (ii) outlining actionable opportunities to improve the industrial data management and integration process. Although this activity is not strictly an engineering analysis of system physics, it is a critical enabler for PHM on legacy equipment, where heterogeneous repositories, incomplete histories, and imperfect traceability are often the dominant practical constraints to robust condition monitoring and prognostics. For example, Figure 5.11 presents an anonymized example of an availability chart.

Correlation Matrices

In the context of the analysis of CF trends, as commonly performed during Exploratory Data Analysis (EDA), a set of correlation matrices for the CFs has been performed to highlight possible duplicates in signals. More in detail, a correlation matrix is a statistical tool utilized to summarize and investigate the relationships among multiple variables within a dataset. Each cell in the matrix contains a correlation coefficient that quantifies the strength and direction of the linear relationship between the two corresponding variables. Correlation matrices are particularly valuable in EDAs, as they help researchers identify patterns, detect dependencies among variables, and make informed decisions for subsequent statistical modeling, such as regression analysis.

The formula for calculating the correlation coefficient between two variables X and Y can be expressed by Equation 5.11

$$\rho_{X,Y} = \frac{\mathrm{Cov}(X,Y)}{\sigma_X \sigma_Y} \tag{5.11}$$

Where:

- $\rho_{X,Y}$ refers to the correlation coefficient between variables X and Y.
- $\mathrm{Cov}(X,Y)$ refers to the covariance between X and Y.
- σ_X indicates the standard deviation of variable X.
- σ_Y indicates the standard deviation of variable Y.

In particular, to analyze the FP temporal evolution identified through the SMs, correlation matrices were constructed at the aircraft level using 50 FPs derived from SM features. Only raw signals were used without including synthetic features. Specifically, for each flight, the first four SM were calculated and combined into a single feature vector representing the flight conditions. These feature vectors formed a dataset for each aircraft. Correlation matrices were then computed across the 50 FPs within this dataset, providing a summary of the dependence structure among FPs at the aircraft level. To evaluate the consistency of these correlation patterns across the entire fleet, the individual aircraft correlation matrices were aggregated by calculating their element-wise mean and standard deviation. This approach enabled the quantification of the average correlation structure and its variability across different aircraft, thereby assessing whether similar trends are observed at the fleet level.

For reasons of brevity, only the correlation matrices obtained from the kurtosis-based FPs are reported in this thesis. In particular, the following matrices are reported: (i) the kurtosis correlation matrix for a single representative aircraft—specifically, the oldest unit in the fleet and the one with the highest accumulated flight hours (Figure 5.12); together with (ii) the element-wise mean correlation matrix across all aircraft (Figure 5.13) and (iii) the corresponding element-wise standard deviation matrix computed across all aircraft (Figure 5.14). Equivalent correlation matrices were computed for all SM descriptors and for all aircraft, but are not included here for the sake of conciseness. From an engineering perspective, these matrices provide a description of the underlying flight-mechanics couplings among variables. Strong (positive or negative) correlations highlight consistent co-variations driven by shared load paths, control actions, and operational regimes. Figure 5.14 shows that the the deviation is overall minimal and that the average matrix can be representative for the whole fleet. In fact, all variance values are very low with the highest (0.31) related to the M_x moment applied to the right HT and the M_x applied to the left HT. This may be justified by the fact that in the lateral-directional plane, these surfaces are often the main ones employed to change the direction of the flight. As a result, a possible conclusion is that despite the very low variability, this correlation is the one which better discerns the different flights. This is justified by the fact that another of the highest value (0.23) is linked to the correlation value between the lateral and longitudinal position of the pilot control stick. If we now take a look at Figure 5.13, where the kurtosis based trends are analyzed, we can see some interesting trends. For instance, the buffeting coefficients are substantially uncorrelated from the other variables together with flap position, pedal position, while the force F_Z and the moment M_x applied on the left HT appear to be strongly correlated with the same signals on the right HT. Other flight mechanics relations are clearly visible: loads and angular speeds are correlated and the loads on the fuselage reflect their influence on the tail moments and forces. For instance, the force F_X on the fuselage is strongly correlated with N_X the moment M_Z and the force F_X acting on the HT. Finally, as expected, the static temperature shows a correlation with the pressure altitude and the Mach shows a strong correlation with the TAS. These analyses were essential for presenting the results to fleet managers and maintenance personnel, facilitating further investigation of the relationships among variables and providing statistical validation for the correlation function.

FH and UR analyses

A specific analysis focused on a study that investigated a possible correlation between aircraft usage in terms of FHs and the number of failures. In this case, only the aircraft TN is considered without explicitly accounting for the swapping and changes of PAS SNs. The results show no evident correlation between sheer usage and the

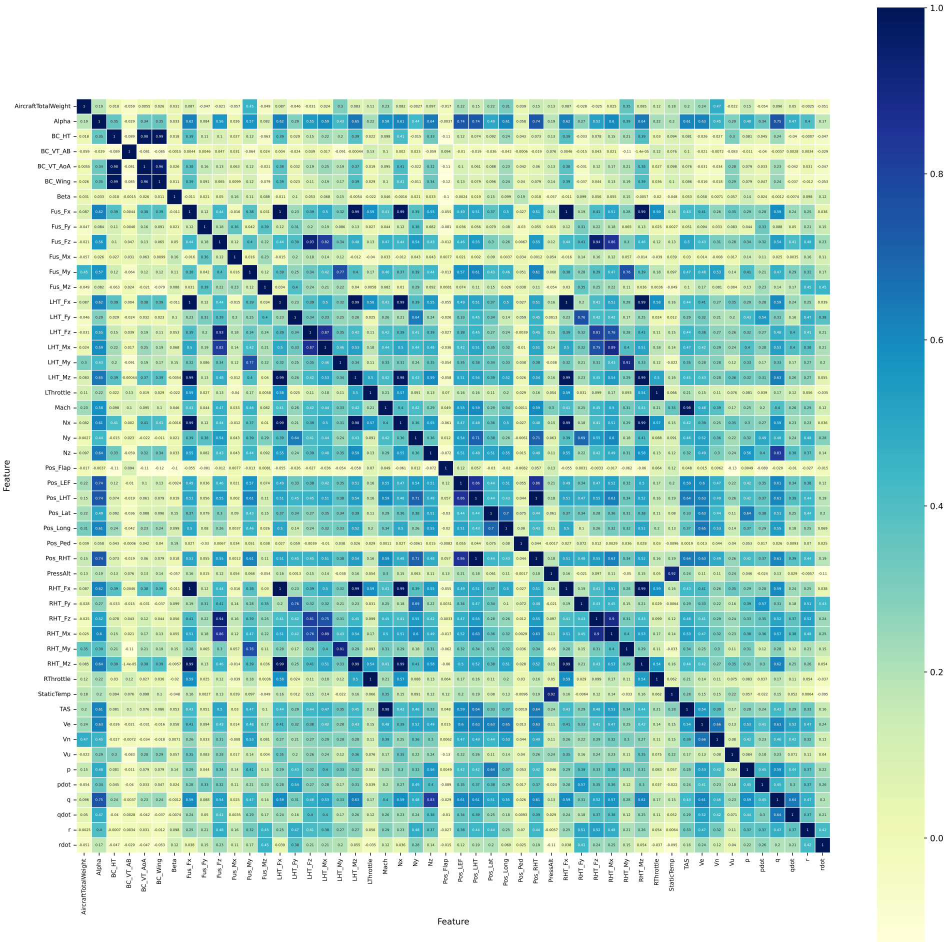

Fig. 5.12 Kurtosis correlation matrix. Correlation matrix of the 50 kurtosis-based FPs computed for the selected aircraft (oldest unit in the fleet and highest accumulated flight hours).

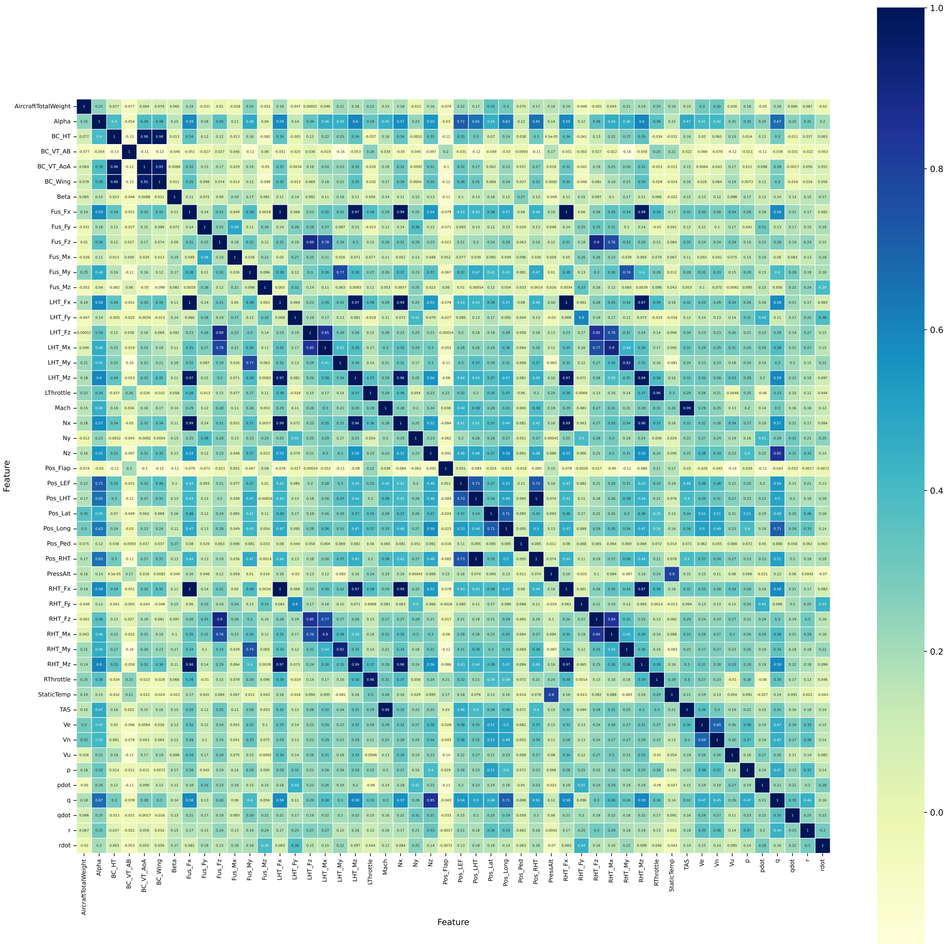

Fig. 5.13 Element-wise mean correlation matrix of the 50 kurtosis-based FPs across all aircraft, summarizing the average dependence structure at fleet level.

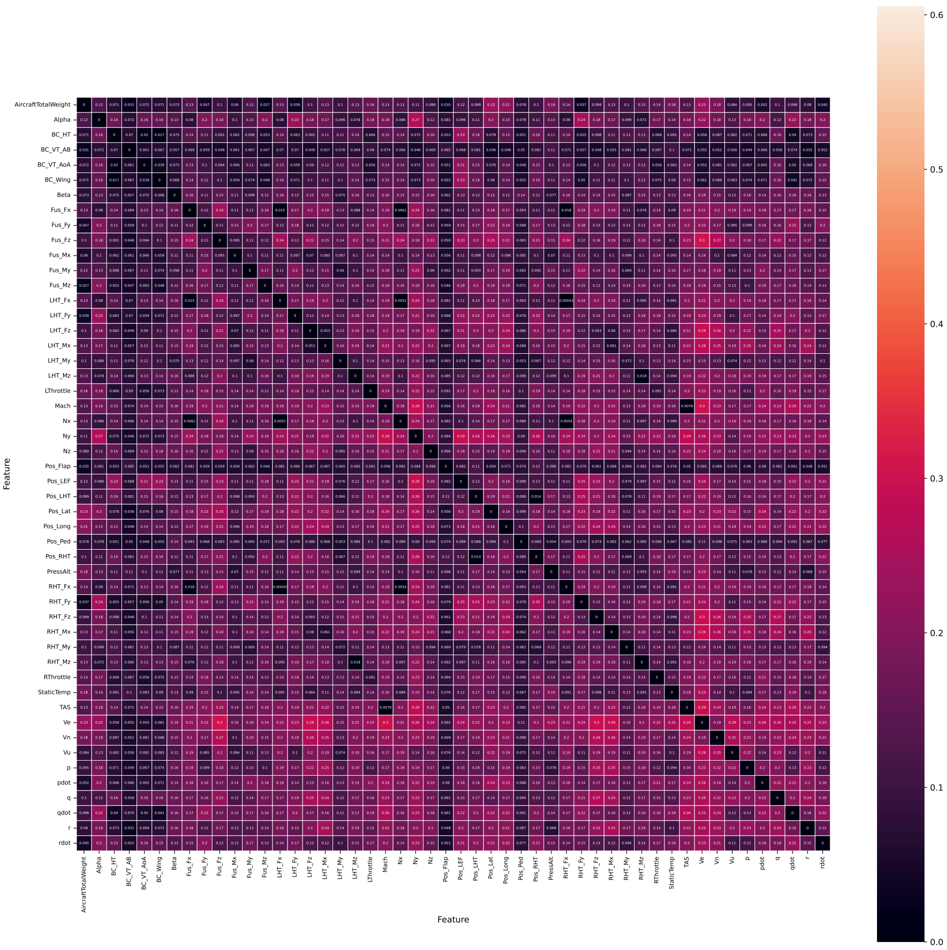

Fig. 5.14 Element-wise standard deviation of the correlation matrices of the 50 kurtosis-based FPs across all aircraft, quantifying inter-aircraft variability in the correlation structure.

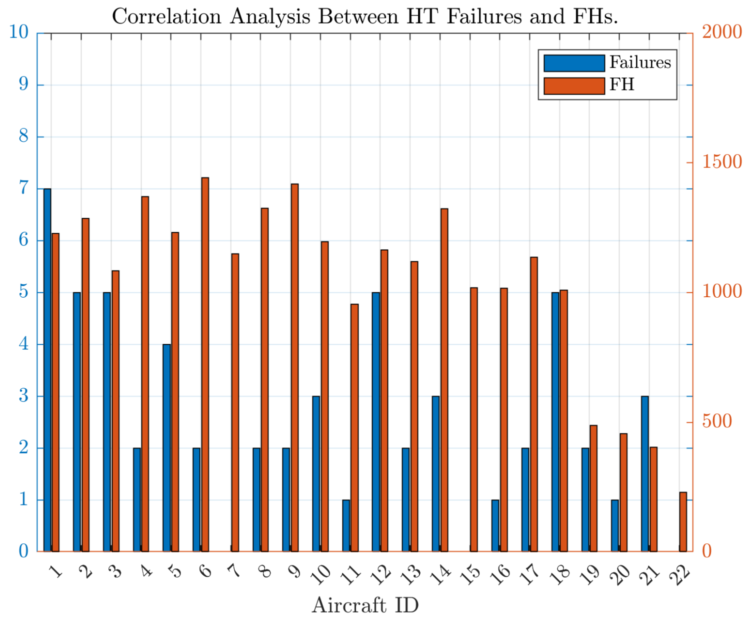

Fig. 5.15 Bar plot of FHs and number of URs for each aircraft.

number of URs related to HT PAS in a specific TN. Figure 5.15 shows a bar graph showing the total FH flown by each TN fleet and the total UR related to the PAS suffered. It is clear that no correlation can be found between the number of failures and the flight hours. An additional statistical analysis is performed by looking at the Pearson's coefficient: considering a standard significance level of 0.05, a low value of 0.314 can be seen with a P-value of 0.154. In conclusion, the number of flight hours is deemed not to be associated with the number of HT-related failures at all.

5.2 Flight Clustering Module

The Flight Clustering Module (FCM), part of the Data Processing, Features & Prognosis Layer, makes use of altitude trends to create a set of flight clusters defined as Clustered Mission Types (CMTs). In this case, it is helpful to review the primary structure of the PHM framework to better understand the interactions of the FCM within the other modules, as illustrated in Figure 5.16. This flowchart provides an alternative but equivalent perspective to the layered overview presented in Figure 3.2, with the FCM highlighted with a thicker frame.

The FCM module aims to group and classify flights according to their mission profiles and identify distinct CMTs that can be used to predict future aircraft usage. The fundamental principles of FCM are reported in a published conference paper [7].

To link actuator prognostics with aircraft use, it is important to accurately describe and quantify mission profiles [85]. An effective approach is to group similar flights into clusters and associate each cluster with a defined probabilistic pattern of expected actuator usage. Once this classification is available, it becomes possible to simulate damage evolution by assembling sequences of missions, whether for short-term planning or long-term projections. As a result, the clustering process becomes a central part of the broader PHM system, providing a connection between operational data and component-level health estimation. A breakdown of the FCM is reported in Figure 5.17, while a detailed view of the Custom Pre-processing block is reported in Figure 5.18. These figures clarify the internal logic of the module and to make explicit how the different inputs are transformed into the final sequence of future CMTs. In particular, Figure 5.17 provides a system-level view of the full workflow, showing how clustering data, simulation horizon, and flight plan are combined through SOM-based processing, CMT extraction, and short-/long-term sequence generation. Figure 5.18, in turn, focuses on the first stage of this workflow, detailing the operations required to convert raw fleet flight data into a standardized time-series dataset suitable for SOM training. These pictures provide a more in-depth analysis of the overall layered PHM framework presented in Figure 3.2, with a focus on the interaction between the FCM and the other modules as a standalone process.

As soon as the Clustering Data is inherited from the Pre-processing module, it is sent to the Custom Pre-processing block which performs additional pre-processing, specifically for clusterization, to the data. First, since clustering is performed on a carefully selected subset of the available dataset, the relevant data subset is selected. This subset captures the diversity of missions without overloading the clustering process with redundant patterns. This subset was obtained from data collected from four specific aircraft, totaling more than 3,000 flights. These aircraft were selected as representative of the overall fleet because they are among the oldest in service, have accumulated the highest number of flights, and have operated across different bases.

Altitude has been selected as the primary parameter for this analytical framework because of three main factors.

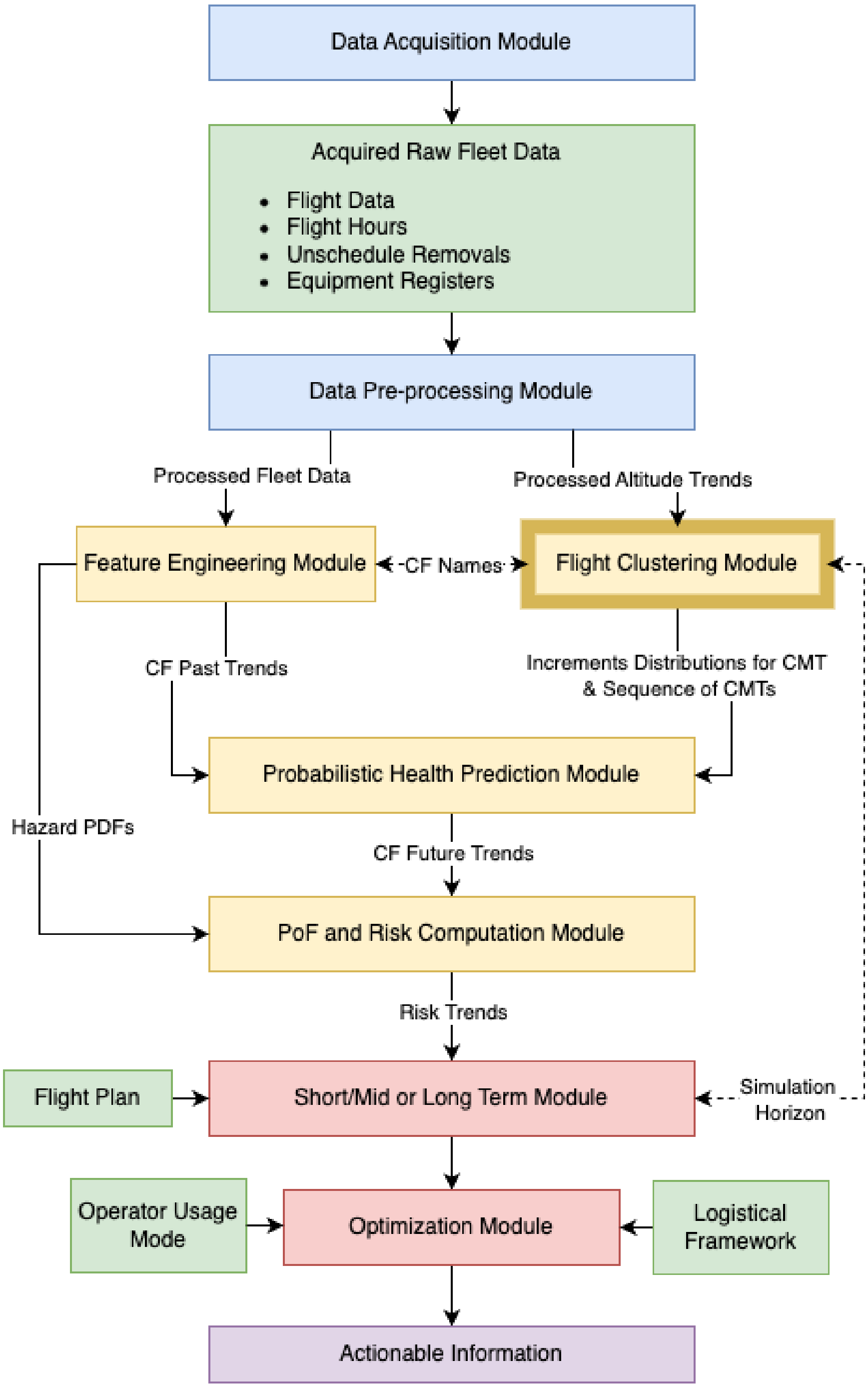

Fig. 5.16 Flowchart of the PHM framework. Each color represents a functional layer: blue represents Layer 1 (Data Management, ETL & Pre-processing), yellow Layer 2 (Data Processing, Features & Prognosis), and red Layer 3 (Health Management, Strategies & Support). The location of the Flight Clustering Module is highlighted by thick edges. Green identifies data sources (internal or external), while purple highlights the framework output.

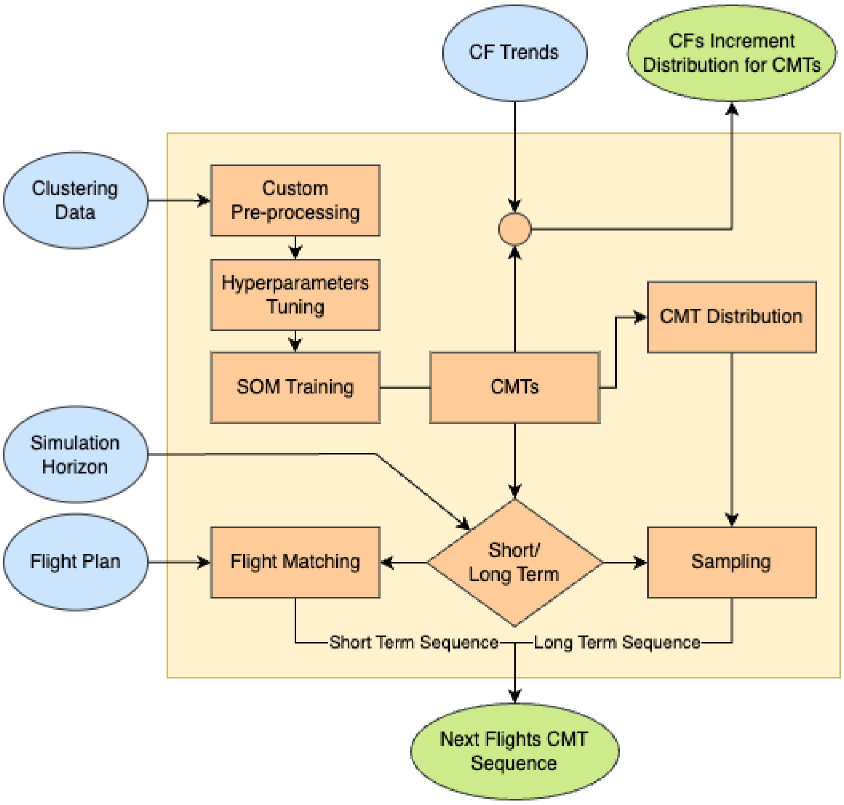

Fig. 5.17 Flowchart of the FCM module. The blue circles represent module inputs whereas the green circles represent the module output.

- First, altitude trends demonstrate a robust correlation with horizontal tail (HT) performance in the longitudinal plane, providing valuable insights into the aircraft's dynamic behavior. Despite the fact that the AJT is equipped with an all-moving HT that merges the functions of the stabilizer and elevator, the main flight mechanics plane which is impacted by its movement is the longitudinal one. Altitude is deemed the best FP to concentrate HT usage in relation to the aircraft macroscopic movements.

- Second, mission profiles are traditionally and consistently represented through altitude-based graphical visualizations, enabling standardized analysis across different flight scenarios.

- Third, altitude-based analysis offers an intuitive and easily comprehensible approach to understanding aircraft performance, making it an accessible metric for various stakeholders involved in aircraft operation and maintenance [125]. This choice of parameter facilitates both technical and practical applicability in monitoring and grouping aircraft operations.

The cluster analysis has been performed on a subset of flights to limit the computational cost. In fact, clustering operations on time-series data are very demanding, and it is deemed that the variety of flights can be represented enough in a smaller subset of flights. Thus, a training subset comprising more than 3,000 flights for a total of four TNs has been selected. A tailored pre-processing technique designed for this analysis has been created that includes resampling reshaping, scaling, and final resampling.

As shown in Figure 5.18, each flight altitude time series has been initially resampled at a uniform rate of 1 Hz using linear interpolation. The resulting data have then been formatted into a custom Python time series structure. Following this, variance-based standardization was applied to normalize the data across the dataset. Finally, each flight has been resampled again at a fixed length of 1000 data points, producing a consistent input shape of 1×1000. This final resampling step is necessary to comply with the input requirements of the SOM algorithm, which requires a uniform time series length across samples. The choice of 1,000 points has been chosen considering the average flight duration (approximately 1 hour) and represents a balance between capturing sufficient temporal detail and maintaining low computational costs. Once preprocessed, the altitude time series

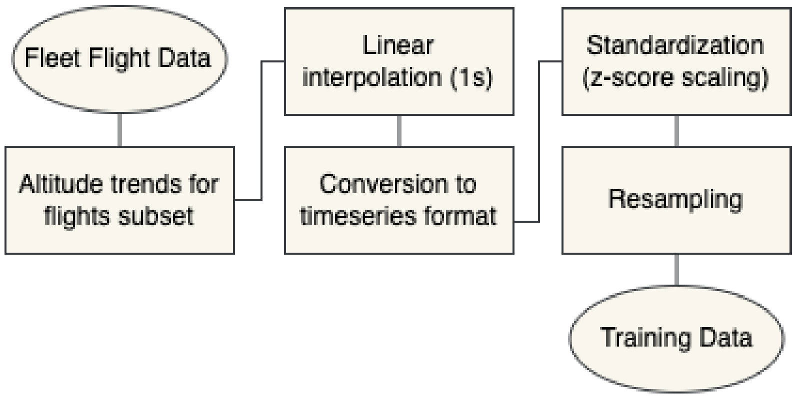

Fig. 5.18 Custom pre-processing flowchart: a detailed view of the "Custom Pre-processing" block of Figure 5.17.

were passed into a SOM model (whose hyperparameters have been previously tuned) specifically adapted for time series clustering to account for the visualization of temporal misalignments. The Python package MiniSom was used in this project [126].

From a data perspective, flight data can be categorized as a multivariate time series dataset whose altitude trend represents a univariate time series. From the task point of view, given the lack of labels identifying the mission performed during flights, unsupervised clustering remains the only viable option. The review of time-series clustering reveals a limited array of options available in the current literature. Recent reviews highlight a clear progression from traditional clustering pipelines to neural and self-supervised approaches but with still limited options. [127, 128]. Among the most commonly utilized methods are time-series k-means [129, 130], k-shape [131], and Self-Organizing Maps (SOMs) [132]. The decision to employ SOMs is primarily motivated by several advantages that they offer, including lower computational costs, ease of implementation, and exceptional interpretability of results due to their powerful ability to reduce data dimensionality. Additionally, deep self-supervised encoders learn task-agnostic embeddings that can be clustered downstream. In the context of the presented PHM setting, which involves unlabeled univariate data prioritizing interpretability, efficiency, and deployability, a SOM offers an optimal solution. SOMs learn a two-dimensional grid that preserves topology, transforming the latent space into an interpretable map of operational clusters. Clusters naturally form as coherent regions in the U-matrix with physically meaningful prototypes that can be examined by engineers. This approach requires no

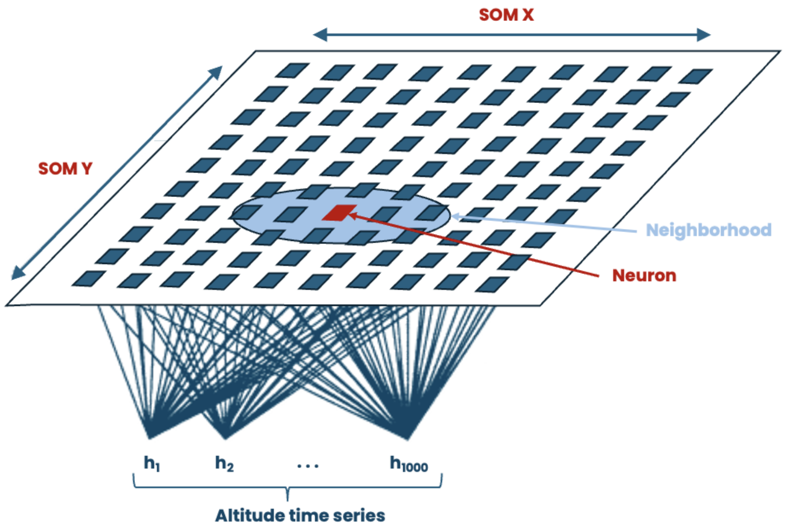

Fig. 5.19 Graphical representation of the SOM 10x10 grid along with the neighborhood space whose weights are impacted by the neuron training.

extensive time alignment and, unlike deep learning architectures that often involve complex hyperparameter tuning and data augmentation, SOMs are computationally efficient, scale effectively to large datasets, and can be trained rapidly on standard hardware. Additionally, SOMs support incremental updates when new data become available, making them particularly suitable for fleet monitoring applications. Let us now consider n as the number of time series ($\simeq 3000$), m as the length of each series (1000), k as the number of clusters (if known), G the number of SOM prototypes and E the number of training epochs. From a computational standpoint, k-Shape requires per-iteration complexity $O\big(\max\{nkm\log m,\ nm^2,\ km^3\}\big)$ [133]. Soft-DTW k-means incurs quadratic cost in the sequence length m, yielding $O(nkm^2)$ per iteration plus barycenter updates [134]. In this context, the Soft-DTW barycenter is the sequence that minimizes the total Soft-DTW distance to all time series assigned to the cluster. In other words, it is the “average” time series under elastic temporal alignment. In contrast, a SOM compares each series with G prototypes and updates neighbors, training in $O(nGmE)$ and inferring in $O(Gm)$ per series, without building full $n \times n$ distance matrices. For instance, when tested over a small subset (250) of the flight databases, the SOM run in 0.24 % of the k-shape running time and comparable time (around 4s) with respect to K-means.

SOMs, also referred to as Kohonen maps, are a type of unsupervised machine learning technique that use competitive learning, rather than error-correction learning techniques, such as backpropagation with gradient descent, to organize data into multi-dimensional maps (Figure 5.19). In other words, they are a type of neural network that establishes a topological relationship between clusters of data. Since their introduction, SOMs have been used across a range of application domains, including biology, geology, healthcare, industry, and geography as an effective and interpretable tool for exploring, clustering, and visualizing high-dimensional data sets. SOMs operate in two distinct phases: training and mapping. In the training phase, a dataset, referred to as the "input space", is processed to construct a lower-dimensional representation of the input data, called the "map space." In the mapping phase, the system classifies new input data based on the established map. Competitive training differs from traditional learning techniques commonly used in AI and involves several key steps that enable the network to learn from input data through a competitive learning process.

1. Initialization. Initially, the weights of the neurons in the SOM are randomly assigned to small values. Each neuron in the output layer corresponds to a weight vector that represents a point in the input space. A seed can be decided so that the test can be repeated accordingly.

2. Training. During this phase, an input sample is selected from the training dataset. The SOM calculates the Euclidean distance between this input training vector and the weight vectors of all neurons on the map. The neuron with the smallest distance to the input vector is identified as the Best Matching Unit (BMU). This process can be mathematically expressed as follows:

$$j = \arg\min_k \|\mathbf{x} - \mathbf{w}_k\| \tag{5.12}$$

   where $\mathbf{x}$ is the input vector, $\mathbf{w}_k$ is the weight vector of neuron k, and j is the index of the BMU.

3. Weights update. After identifying the BMU, the SOM determines which neighboring neurons should also be updated. This is done using a neighborhood function that defines how far from the BMU other neurons can be affected by the learning process. The influence decreases with distance from the BMU.

In this step, both the BMU and its neighboring neurons adjust their weights to become more similar to the input vector. The weight update rule is given by:

$$\mathbf{w}_j(t+1) = \mathbf{w}_j(t) + \lambda(t) \cdot h_{j,k} \cdot (\mathbf{x} - \mathbf{w}_j(t)) \tag{5.13}$$

where: t is the current iteration, $\lambda(t)$ is the learning rate at time t, $h_{j,k}$ is the neighborhood function that determines how much influence the neuron j has based on its distance from the BMU (in this case a Gaussian function), $\mathbf{x}$ is the input vector and $\mathbf{w}_j(t)$ is the weight vector of the neuron j.

4. Iterations. The above steps are repeated for a predetermined number of iterations.

SOMs consist of a network composed of two layers: an input layer and an output layer of interconnected nodes, commonly referred to as neurons or units. Typically, the topology of the output layer is designed as a two-dimensional grid, allowing for straightforward visualization. This ability to visualize the data makes SOM a particularly interpretable clustering method, and this method allows for an intuitive visualization of complex data structures. In fact, in this case, each flight is assigned to the best position in the map (called Best Matching Unit - BMU) and this can be easily visualized in a grid. Once the SOM hyperparameters, such as map dimensions, sigma, and learning rate, are determined, the map is trained using a subset of altitude trends from the dataset. To quantify the clustering results and select appropriate network hyperparameters, classical SOM metrics are used. The architecture of SOMs facilitates a one-to-one correspondence between map cells and clusters, resulting in the identification of distinct clusters that correspond to individual cells within the map. In particular, SOMs can be employed to cluster data in two main ways. On the one hand, one can assign each grid position to a cluster and hence set the number of clusters a-priori based on SOM performance metrics. On the other hand, clustering can be performed even after, by looking at how the SOM has been trained with the help of the U-Matrix (Unified Distance Matrix). This special representation visualizes the distances between neurons (or nodes) in the SOM, providing valuable insights into the structure and clustering of the data represented by the map. The U-Matrix is created by calculating the distances between adjacent neurons in the SOM. These distances are then visually represented, typically utilizing a color gradient in which:

- Dark colors highlight the distances between adjacent neurons, suggesting a significant gap in the underlying data. This may be interpreted as a boundary between clusters.
- Light colors underscore that the neurons are close to each other, suggesting that they represent similar data points and can be categorized within the same cluster.

The second level clustering can be applied on the U-matrix with a wide range of clustering algorithms such as K-means, hierarchical clustering, DB-Scan, etc. This visualization facilitates an intuitive understanding of how data points are grouped within the high-dimensional input space, even without prior knowledge of the clusters. The U-Matrix effectively delineates clusters and their separations, simplifying the process of identifying patterns within complex datasets.

If the advantages of employing SOMs are clear and are mainly related to the intuitiveness of the results, some difficulties arises in the way the performance of SOM results are assessed. In fact, most of the time the training results cannot be compared with a labeled data set. For these reasons, a wide range of metrics have been created to try and provide some indications on the hyperparameter tuning as well as on the overall results. The metrics are often divided into internal and external metrics, with the former using only the data properties to calculate the metric and the latter, which is related to the data labels (when available). The most common internal metrics are the topographic error, the quantization error, the combined error, and the topographic product. The quantization error measures the average distance between each input vector and its BMU on the map. It represents how accurately the SOM weight vectors match the input data patterns. The lower quantization error indicates a better representation of the data by the SOM. The topographic error evaluates how well the SOM preserves the topological structure of the input data during the mapping process. It specifically measures the degree to which the relationships between data points in the input space are maintained in the output space represented by the SOM. The topographic error is defined as the average geometric distance between the BMU and the second-best matching unit for each input vector. If these two units are neighbors on the map, it indicates that the topology has been preserved for that input. Topographic error is essential for evaluating local discontinuities in mapping, which can impact how well a SOM captures complex patterns in high-dimensional data. The combined error is a way to merge the information provided by the quantization and topographic error. The topographic product assesses the preservation of neighborhood relationships between the input space and the resulting map. This measurement relies solely on the prototype vectors and the topology of the map. It can indicate whether the dimensionality of the map is suitable for adequately representing the dataset or if it has introduced neighborhood violations due to map distortions. As a result, it can be used to guide the aspect ratio of the output map, since it serves as a criterion for optimizing the dimensionality and structure of the SOM output space.

Despite the popularity of applications leveraging SOM high dimensionality reduction capabilities, the applications involving time series analysis are definitely scarcer. As a result, a hyperparameter tuning pipeline had to be designed. The first part of the tuning focused primarily on two key parameters: the total number of clusters and the aspect ratio of the SOM grid. Although common heuristics suggest estimating the number of clusters as $5 \cdot \sqrt{NumberofInputs}$, this formulation produced an impractical high number for operational use (216). We set the number of clusters to 100 after trial-and-error tuning to balance interpretability and resolution, incorporating feedback from operations personnel. The map aspect ratio, defined as the ratio between the number of neurons along the horizontal and vertical dimensions of the SOM grid, was optimized by evaluating the topographic product, a metric that quantifies how well neighborhood relationships in the input space are preserved after projection onto the map. In general, values closer to zero indicate a better preservation of topology, whereas larger negative values indicate a poorer mapping quality. This metric has been adopted in several SOM optimization studies to guide the selection of a suitable grid geometry [135]. Figure 5.20 shows the mean topographic product obtained for different SOM shapes having the same total number of neurons but different aspect ratios, ranging from highly elongated grids to a square grid. On the x-axis, each label reports both the aspect ratio and the corresponding grid shape; for example, 100 corresponds to a 100 by 1 map, whereas 1.0 corresponds to a 10 by 10 square map. The red line represents the mean topographic product over 10 independent trainings for each configuration, while the blue shaded band represents the associated standard deviation, thus providing an indication of the variability due to random initialization. The results show a clear improvement as the map becomes less elongated and more balanced, with the topographic product progressively approaching zero. The best result is obtained for aspect ratio 1.0, corresponding to the 10 by 10 square grid, which therefore provides the most suitable

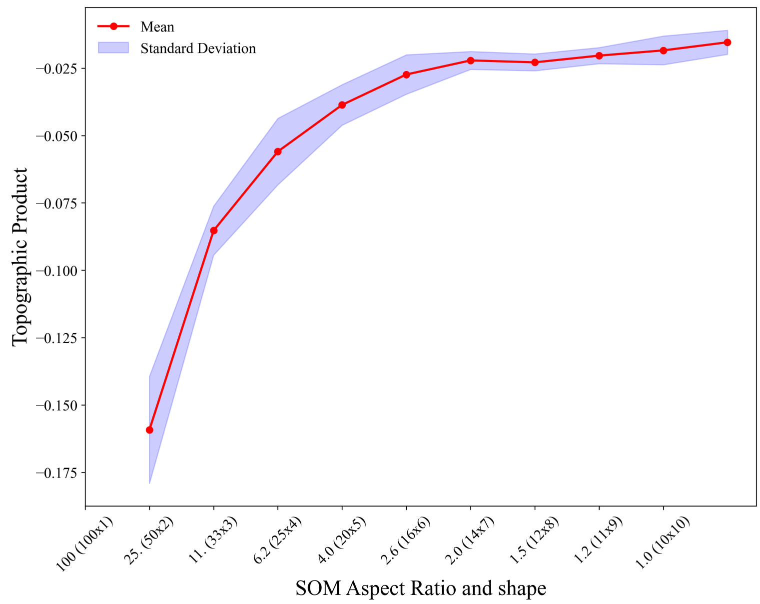

Fig. 5.20 Topographic product trend in function of the aspect ratio along with the standard deviation band. In this instance, the SOM was trained considering the other parameters as default values.

geometry in terms of topology preservation. To ensure statistical consistency and reduce the effect of random initialization, each SOM configuration with fixed aspect ratio was trained independently 10 times. The next step has involved the selection of the two main hyperparameters: sigma and the learning rate. In particular, an end-to-end hyperparameter tuning pipeline based on grid search is proposed as reported in the pseudocode 1, exploring a total of 504 parameter combinations. The scoring **S** is defined in Eq. 5.14. After the selection of these last parameters, the network topology is fixed and the SOM architecture is ready to be trained and employed in the FCM. In fact, as already stated, SOMs operate in two distinct phases: training and mapping. In the training phase, the input dataset is used to build a lower-dimensional representation.

$$\mathbf{S} = QE + TE + 2 \cdot (1 - \mathbf{U_{rate}}) \tag{5.14}$$

where QE is the quantization error, TE is the topographical error, and $\mathbf{U_{rate}}$ represents the utilization rate (i.e., the percentage of neurons that are the BMU for at least one flight). The tuning operation has been performed with a fixed seed to safeguard reproducibility. The solution space is reported in the tridimensional graph in Figure

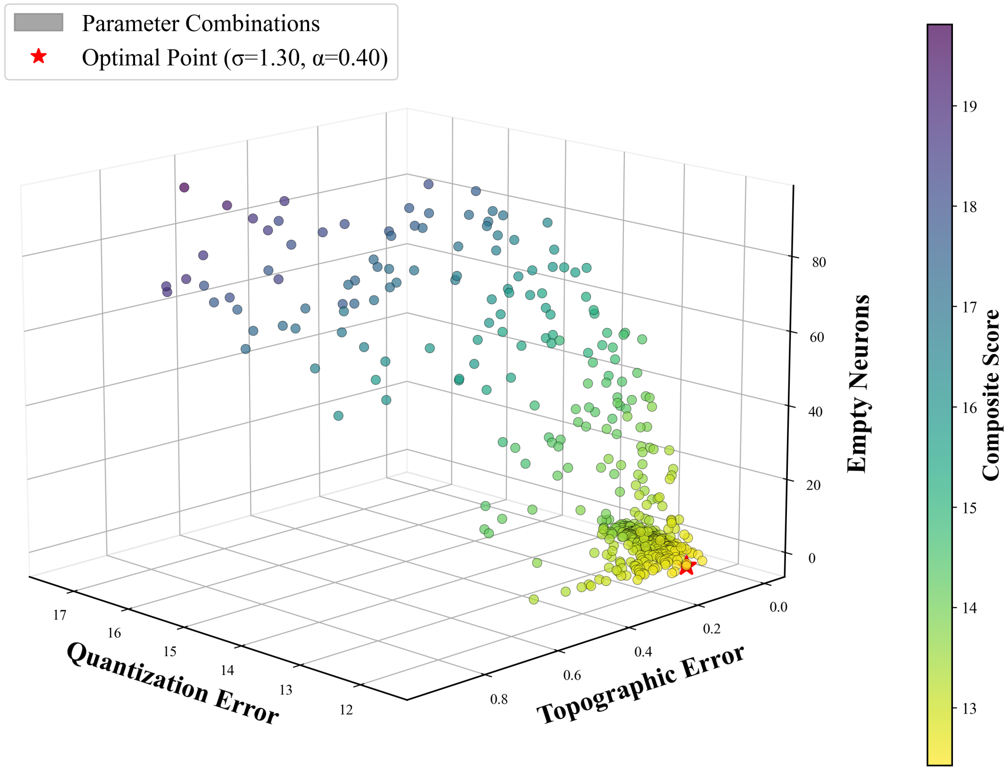

Fig. 5.21 SOM hyperparameters 3D space: quantization error, topographical error and utilization rate. Grid search has been used to identify the best parameters.

5.21: the final values of sigma and learning rate are 1.3 and 0.4, respectively. The pseudo code to perform hyperparameter tuning is reported in the Algorithm 1 in Appendix.

During this phase, the SOM learns to organize the input data by adjusting the weights of its neurons based on their similarity to the input vectors. Each input is compared to all neurons on the map, and the one with the closest BMU is identified. The BMU and its neighboring neurons are then updated to more closely resemble the input, gradually shaping the map to reflect the underlying structure of the dataset, in a procedure called competitive training. Throughout training, various classical SOM metrics are used to evaluate the quality of the clustering and guide the tuning process. The training lasts 100 epochs as already after 50 the errors stabilize and further epochs do not increase the accuracy of the model. Once training is complete, the mapping phase begins. In this phase, new input data can be projected onto the trained map. Each new input is assigned to its BMU, effectively classifying it based on the structure learned during training. This process enables intuitive visualization and interpretation of complex, high-dimensional data by preserving topological relationships, data points that are similar in the input space remain close together in the map space, as better explained in the next paragraph.

The high explainability of SOMs is clearly illustrated by the joint interpretation of Figure 5.22 and Figure 5.23. In both figures, the same (10 BY 10) SOM grid is shown, where each square corresponds to one neuron of the map, and therefore to one prototype flight profile learned during training. In Figure 5.22, each CMT is associated with its winning neuron, that is, the neuron whose weight vector best matches that specific flight. This allows each flight to be positioned on the map according to its similarity with the learned prototypes. Figure 5.23, in turn, provides the complementary information needed to interpret this arrangement: the U-matrix reports the Euclidean distances between neighboring neurons in the weight space, thereby highlighting where the map is locally homogeneous and where more marked transitions occur, like a geographical map with hills and valleys. Low distances indicate regions of similar neurons, whereas high distances identify boundaries between different groups of flight profiles. For this reason, the two representations are directly related and mutually complementary: the timeseries map shows where the flights are placed, while the U-matrix explains how close or separated the underlying neuron prototypes are. Together, they make it possible to visually identify clusters of similar CMTs, understand transitions between different flight behaviors, and verify that the optimized SOM training has organized similar nodes in neighboring regions of the grid. The grid is now analyzed more in details.

As a lead-in fighter training platform, the AJT operates within a structured syllabus of recurring mission types, yet individual flights exhibit substantial variability even within identical mission categories. A comprehensive analysis should incorporate both Figure 5.22 and Figure 5.23, as the information they present is interrelated. For instance, the U-matrix highlights a mega cluster of similar CMTs in the center of the grid; this indication is backed up by Figure 5.23 which tells us that that mega cluster represents common cruise flights.

It is challenging from the mission profiles displayed in Figure 5.22 to associate typical labeled mission. The AJT under investigation is a lead-in to fighter training aircraft which therefore is characterized by recurring mission agenda, yet the flight

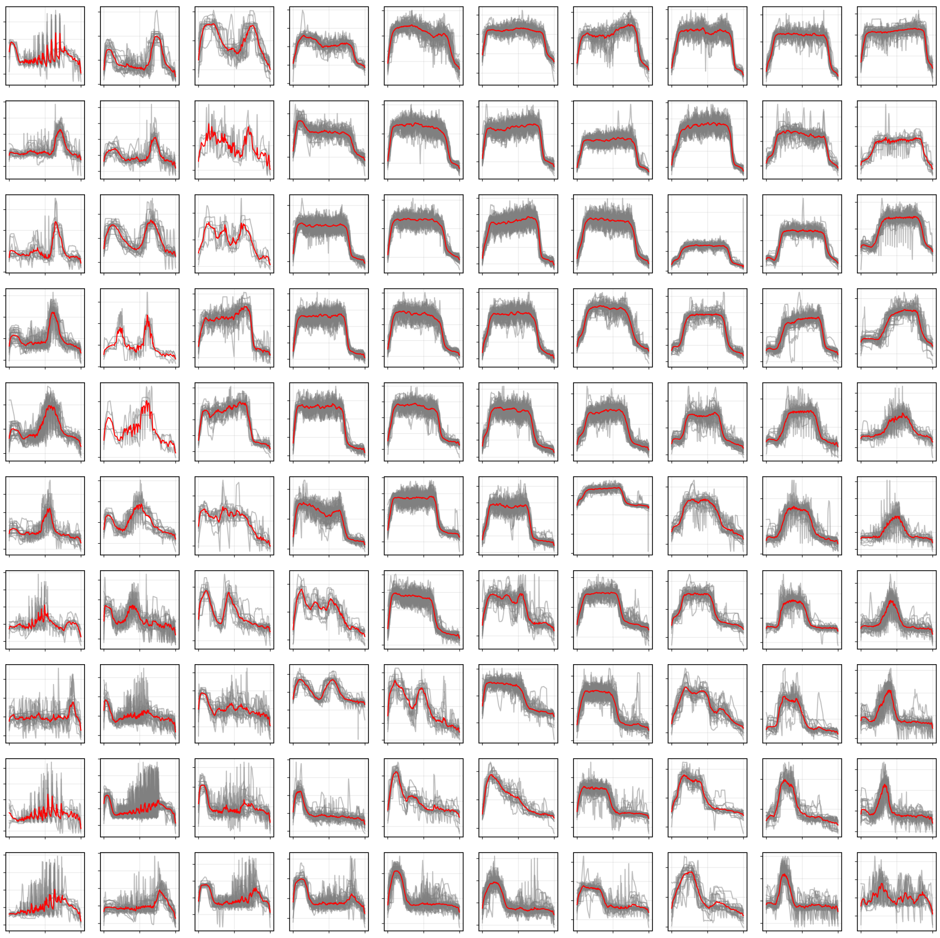

Fig. 5.22 CMT clusters with associated altitude trends. The gray lines represent the combined trends within each CMT, while the red line indicates the designated trend for that specific CMT.

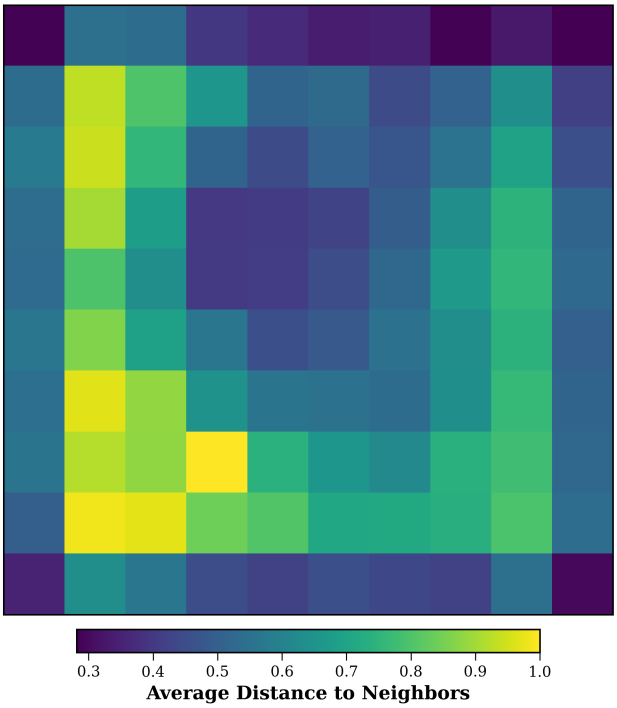

Fig. 5.23 U-matrix for the trained SOM: dark regions indicate similar clusters while lighter regions indicate non-similar clusters. The average number of flights per cluster is 32.5.

trends can vary drastically between flights of the same mission type. Nevertheless, general considerations can be made. Steady high-altitude trends could be representative of basic navigation missions, but could also be indicative of dog fighting scenarios. Lower altitude trends before landing could likely be associated with approach missions. Low-altitude trends could be lastly representative of air-to-surface situations. Finally, several altitude variations during single flights could be indicative of formation flights or otherwise tactical missions. It should be mentioned at last that single flights often combine several of the many mission types that make up a full pilot training syllabus.

It is important to note that the observed similarities among certain CMTs should not be interpreted as errors in data processing or clustering methodology. Rather, these similarities reflect the effectiveness of the optimized 10-by-10 network topology established during the hyper-parameter tuning phase. This approach has demonstrated that these flight profiles are indeed common within the dataset, featuring multiple CMTs with only minor variations in their characteristics.

The choice of employing SOMs as a tool to cluster flight missions has been justified by the lack to labeled flight at hand. It is clear that, if a label could be assigned to each flight beforehand, completely different architectures could be employed. However, since the available dataset is completely unlabeled, it is deemed that SOMs are the best solution for this problem.

Common internal metrics include topographic error (0.033) [136], quantization error (11.827), combined error [137], silhouette index (0.047), Davies-Bouldin index (2.390593), and topographic product (0.1023) [138]. A quick implementation of these metrics is provided in a Python package developed by [135]. The excellent topographic error and product indicate strong topology preservation, demonstrating that the SOM effectively maintains neighborhood relationships from the original high-dimensional space to the two-dimensional map. Conversely, the low silhouette and Davies-Bouldin indices suggest room for improvement, highlighting the challenges in identifying well-separated, compact clusters due to the complexity of the data.

The creation of a SOM model is central for the next steps of the framework for many reasons.

- CF Future Evolution. As indicated in Section 5.1, each flight is characterized by a specific set of five CF increments (one for each CF) that outlines the impact of that flight on CF trends. After assigning a CMT to each flight, we can cluster the CF increments of flights that share the same CMT label. This grouping allows us to create distinct distributions of increments tailored to each CMT. Each CMT has a total of 5 distributions of increments (one for each CF). In essence, by aggregating flight increments according to their respective CMTs, we generate a set of five distributions that inherit the unique characteristics associated with that type of sortie. Consequently, each CMT is defined by a collection of five increment distributions that illustrate how it influences the growth of the CFs.

- CMT Historical Distribution. The assignment of a CMT to each flight within the training dataset facilitates the establishment of a statistical distribution regarding the occurrence of CMTs. This distribution effectively identifies which CMTs are more frequently observed and which are less common. The historical distribution is illustrated in Figure 5.24 and can subsequently be utilized to simulate a sequence of long-term future CMTs by sampling from the established historical occurrence distribution.

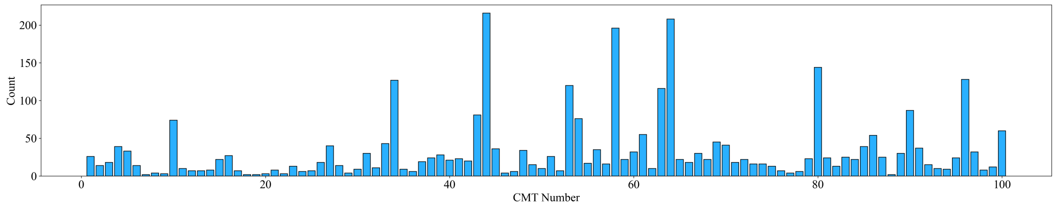

Fig. 5.24 Histogram representing the occurrences of flights for each CMT. This distribution offers a solid way to simulate long-term flight mission sequences by sampling CMTs.

- A-posteriori Classification. It is now possible to use the trained SOM as a classifier so that every new flight can be classified and linked to a specific CMT.

The data obtained can be used to create a sequence of CMTs for the next flights. The sequence creation can be created based on two different requirements: short-to-mid term and long term prediction. The short-to-mid term approach can be used when a fleet operator wants to assess the health status of the PAS for a set of known future missions in the near future. In fact, once the specific missions for a specific aircraft TN are known, these can be translated into CMTs and then assembled into a sequence of CMTs for the next flights. On the other hand, a long term prediction can be performed to assess the condition of the PAS without the need for a predefined set of CMTs. In this case, the sequence of CMTs can be sampled based on the CMT occurrence distribution and can be as long as desired, involving a wide range of flights. The decision regarding which strategy to employ is made within the Health Management, Strategies & Support Layer, which is why a dotted arrow connects the two layers. Additional details on both short-to-mid and long-term simulations can be found in Section 6, where the third layer is discussed in depth.

Figure 5.25 shows the process of CMT sequence assembly. The short to mid term approach requires the information from a custom flight plan based on next flight missions and provides more precise results while sampling from the historical CMT distribution can provide long term predictions. Once the CMTs have been identified, the sequence can be assembled.

Once the sequence of next flight CMTs of length $N_F \in \mathcal{N}$ has been assembled, the Probabilistic Health Prediction Module can take the sequence as input and produce a probabilistic estimation of the CFs trends given the elaborated distributions.

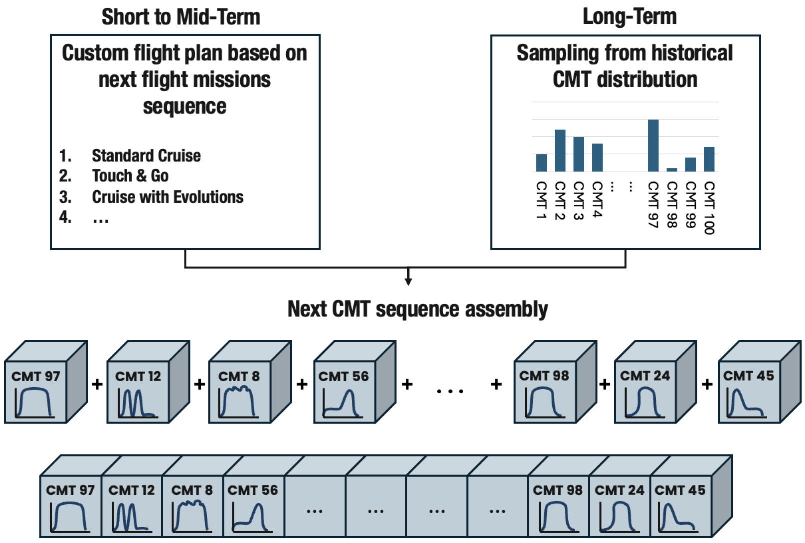

Fig. 5.25 CMT sequence assembly process. The CMTs can be obtained from a known set of next missions (short-to-mid term) or sampled from the CMT distribution for a long term prediction.

In conclusion, the Flight Clustering Module outputs the CFs increment distributions for each CMTs and the sequence of the next flights CMTs.

5.3 Probabilistic Health Predictions Module

The Probabilistic Health Prediction Module serves as the core of the PHM framework, projecting CFs into the future based on the distributions of CF increments and the sequence of upcoming CMTs.

To illustrate, let us consider a specific CF, such as the skewness of the right horizontal tail work; this methodology can be similarly applied to other CFs without any modifications. Although the framework is not explicitly built upon PF techniques, the concept of particles can be effectively employed in this thesis to enhance clarity. Particles are defined as discrete samples that represent potential states of the system; in the context of this thesis, they represent the trends of the CFs. The actions performed in this module can be broadly divided into two categories as follows:

1. Particle Creation

2. Particle Projection

Particle Creation

The initial step involves the creation of the initial set of particles. Due to the limited prior information regarding the system state, particles are generated on the basis of Bayes' theorem.

Bayes' theorem is commonly considered as the foundation of Bayesian statistics, providing a mathematical formulation aimed at the estimation of the probability of a state based on prior knowledge of the system (i.e., the prior distribution) and the likelihood distribution. In other words, Bayes' Theorem provides a mathematical framework for updating the probability of a hypothesis based on new evidence. It is expressed as follows:

$$P(A|B) = \frac{P(B|A) \cdot P(A)}{P(B)} \tag{5.15}$$

where $P(A|B)$ is the conditional probability of the event A given that event B has occurred, $P(B|A)$ is the probability of observing the event B given that A is true, $P(A)$ is the prior probability of A, and $P(B)$ is the total probability of event B. The prior distribution is usually determined from a thorough understanding of the physical system and may assume a Gaussian shape or another custom shape if relevant insights are available.

Prior Distribution In this analysis, the prior distribution reflects the increments of the CFs across all PASs. The rationale for considering data from all PASs is that some PASs may have more extensive historical records than others. Analyzing increments associated only with a specific PAS may enhance precision for that particular system, but could ultimately compromise the overall accuracy of the framework. In fact, some PASs may have limited historical data, and generating particles based on a probability distribution derived from scarce samples can lead to inaccuracies.

This is particularly critical since the generation of particles significantly influences the overall effectiveness of the framework; the initial particles represent the starting points of CF trends, and inaccuracies at this stage may result in persistent errors during the simulation phases. The prior distribution is fitted using a Student's t-distribution, which is selected as the most appropriate fit for this context after a fitting test with Distfit Python package and using RSS as the metric.

Likelihood Distribution Within the framework of Bayes' theorem, likelihood refers to the conditional probability of the observed evidence given a specific hypothesis. It plays a crucial role in updating prior beliefs regarding the hypothesis in light of new evidence.

In this case, the likelihood is determined by the increment distribution of the first CMT in the sequence, as it is known that the CMT of the first flight in the upcoming sequence is available. Using these distributions and applying Bayes' theorem, the initial set of particles is generated, consisting of 1,000 particles, denoted as $N_P \in \mathcal{N}$.

Particle Projection

Each particle position is subsequently updated for each flight according to the increments guided by the CMT sequence projected into the future. The resulting trace, formed by connecting particle positions over time, illustrates a potential evolution of the CF. This process occurs in parallel for each particle, resulting in a total of 1,000 possible traces that provide a statistical basis for predicting CF evolution. These traces, each representing a possible trend of the CF over time, are constructed following the CMT sequence via Monte Carlo simulations.

Monte Carlo simulations are a mathematical technique commonly used to forecast a range of possible outcomes for uncertain events through repeated random sampling. Once the particles are generated through Bayes' theorem, the process is iterated for each new flight in the forthcoming flight sequence.

A new flight is characterized by its CMT, with each CMT linked to a specific increment distribution. The process involves sampling a value from the CMT increment distribution for each particle trace and repeating this operation until the completion of the CMT sequence. This procedure can also be repeated for each CF. Figure 5.26 offers a graphical representation of the particle projection process. In particular, the assembled CMT sequence is inherited from the Flight Clustering Module as already represented in Figure 5.25 and is here used as the basis for the probabilistic projection. According to the increment distribution associated to the CMT, an increment is sampled from the specific distribution (in red). This is the CF

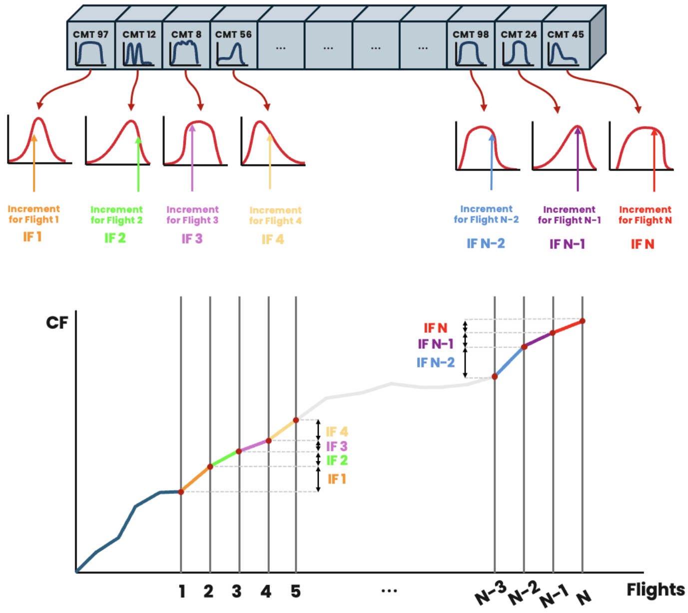

Fig. 5.26 Particle projection process. Each CMT is identified by its altitude trend as reported in the blue plot inside the cube. Each CMT is linked to its red increments PDF, from which a value is sampled.

increment for the next flight. This process is carried out for all CMTs in the sequence, creating a CF trend. In other words, Monte-Carlo simulations are performed for a total of N_F times.

This process is carried out for each particle created during the creation process, as shown in Figure 5.27 and Figure 5.28 where two distinct traces (denoted by different colors) have been created following the explained process. Note that the CMT sequence remains unchanged, as well as the increment distributions. What changes is the sampled increment from the increment distributions.

The sole output of the Probabilistic Health Prediction Module is represented by the future traces of the CFs, which will be further elaborated to obtain PoF and risk indications. Figures 5.29a and 5.29b show the simulations carried out for PAS SN 9, involving 1000 particles and 330 CMT simulation in the future.

In particular Figure 5.29a reports the single traces, whereas Figure 5.29b shows the percentiles of the MC simulations along with the ground truth signal, highlighting an excellent prediction ability even in the mid-term. More details along with the methodology evaluation are reported in Section 5.4.1. Figure 5.30 shows the simulation process at four different values of k_p. More specifically, Figure 5.30 is related

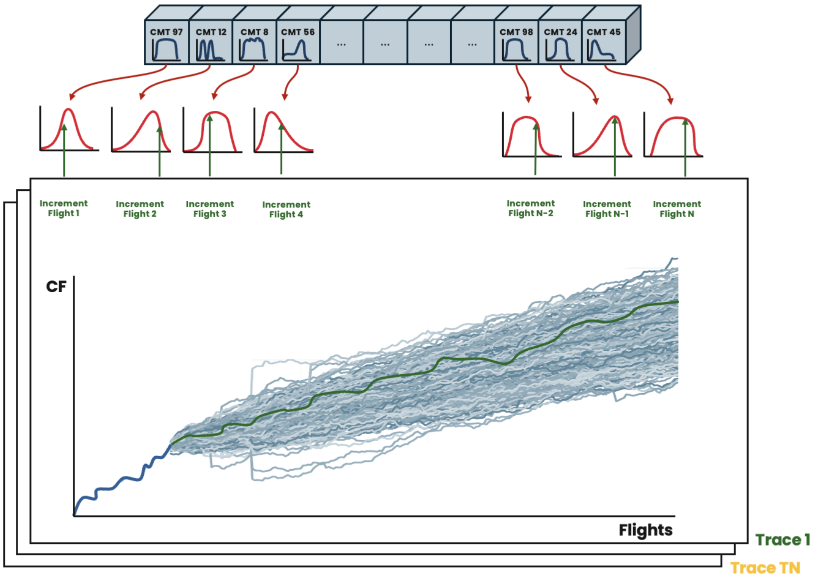

Fig. 5.27 Trace creation process for trace 1.

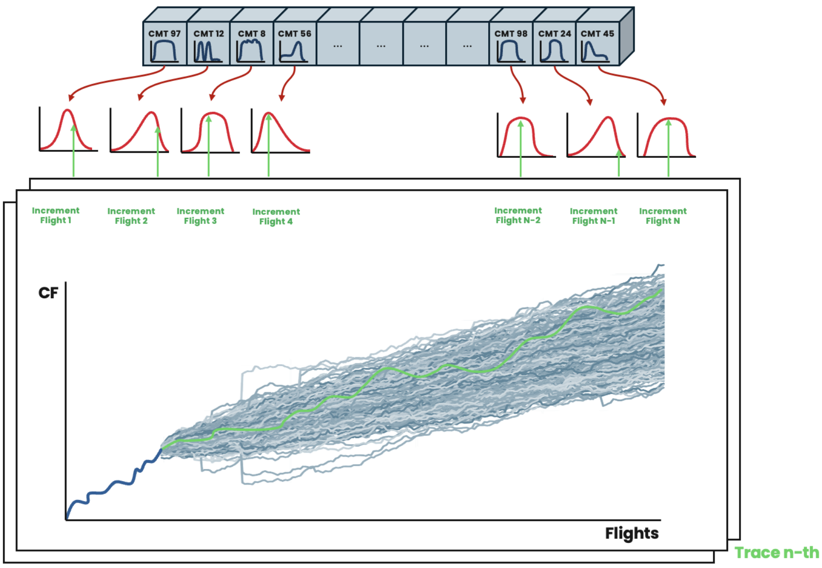

Fig. 5.28 Trace creation process for trace n-th.

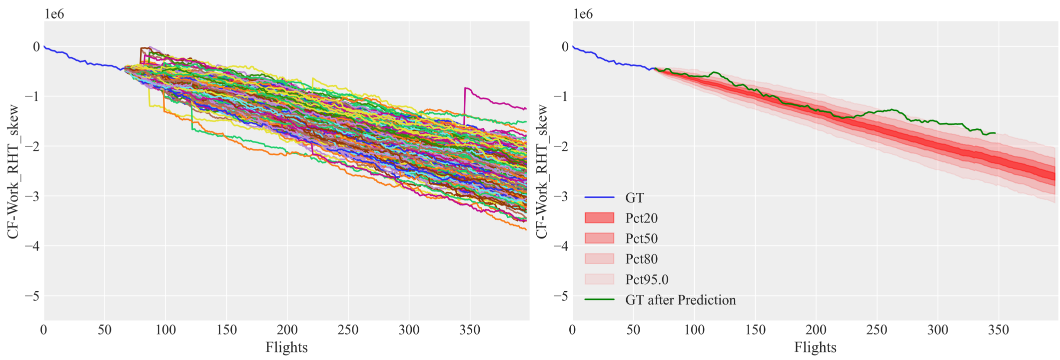

(a) MC traces whose particles have been initially generated via Bayes theorem.

(b) Comparison between simulations and ground truth.

Fig. 5.29 Examples of CF projections in the future through Monte-Carlo simulations for PAS SN 9.

to a different SN than the one considered before to demonstrate the effectiveness of the framework across multiple SNs, as better shown in Section 5.4.1 where the validation of the strategy is reported.

This module's forward-looking design enables a simple transition from a short- to mid-term approach to a long-term perspective, as the only adjustment required consists in the change of the next flight CMT sequence; the health prediction routine itself remains consistent.

5.4 PoF and Risk Computation Module

The Probability of Failure (PoF) and Risk Computation Module utilizes CF traces projections, as provided by the Probabilistic Health Prediction Module, to analyze CFs trends in relation to the hazard threshold distribution derived from the Feature Engineering Module. Thanks to the comparison between the PDF of the system state in a time instant and the PDF of the hazard threshold, it is possible to calculate the PoF and the risk that an UR happens at that time instant. There are multiple formulations of the PoF in literature thanks to the different shape of the hazard threshold. In fact, the hazard threshold may be deterministic (i.e. hard hazard zone characterized by a single value) or probabilistic (i.e. soft hazard zone characterized by a PDF), also called Uncertain Event Likelihood Functions (UELFs) [122].

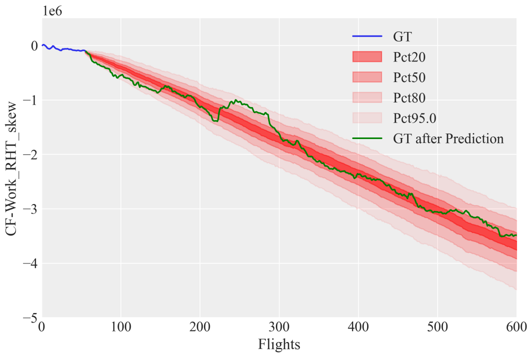

(a) Simulation with k_p set to 56.

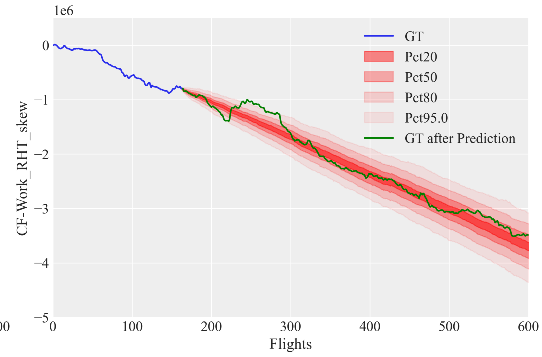

(b) Simulation with k_p set to 164.

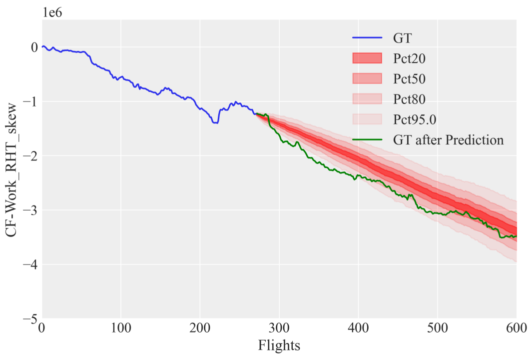

(c) Simulation with k_p set to 272.

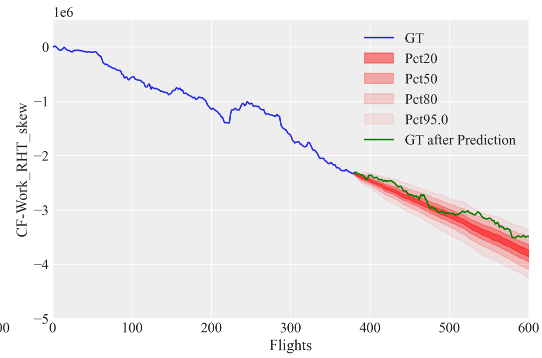

(d) Simulation with k_p set to 380.

Fig. 5.30 MC simulations and results for PAS SN 56 with 1000 particles at four evenly spaced values of k_p with a final flight set to 600. The GT evolution is inside the prediction cone, marked by the 95 percentile.

A hard threshold, which can be re conducted to a deterministic threshold, is a fixed, single value that delineates the boundary of the hazard zone. This approach is characterized by its deterministic nature, where the threshold is a precise, unchanging value. The assessment of one of the traces of the system state relative to the threshold results in a binary outcome.

In contrast, a soft threshold is characterized by a PDF (i.e., UELF), indicating that the hazard zone encompasses a range of values with associated probabilities [139]. This approach reflects the inherent uncertainty and variability in the hazard zone as well as epistemic and systematic uncertainties. The likelihood of the system state being faulty increases progressively as it approaches and exceeds the threshold.

The hard threshold approach, even if simpler to implement and interpret due to its binary nature, it is hardly applicable in real-life operations where greater flexibility has to be taken into consideration. Moreover, as the system condition evolves over time, several approaches can be used to compare the two PDFs by modeling the distribution of the event time; representative examples include first-passage time formulations [140–143] and first-hitting time models [144, 145].

For this purpose, let us assume that a prognostic analysis to predict the next UR is performed at a time k_p. The event $\mathscr{E}$ of interest corresponds to a PAS UR. As already mentioned in the previous paragraphs, the two elements to consider are: the representation of the state of the system in the future flights and the threshold.

In this case, the former is provided by Monte-Carlo simulations of CF trends. This representation corresponds to the stochastic process $X_{k k \in \mathbb{N} \cup 0}$, where the variable k is a time index that takes values in the set of natural numbers (zero included) and indicates the flight number. As common in the evolution of real-world dynamic systems, the process approached in this study is characterized by Markovian proprieties, as the future state is only dependent on the present state. In fact, the CF evolves in time depending on the next CMT which is linked to the increment PDF. The system can then be modeled according to Equation 5.16

$$x_{k+1} = f(x_k, CMT_{k+1}) \tag{5.16}$$

where x_k is the present state and CMT_{k+1} is the CMT of the next flight.

On the other hand, the UELF can be defined as $\mathbb{P}(E_k = \mathscr{E}|x)$ and can be obtained as the cumulative density function of the PDF threshold according to the Feature Engineering Module. It outlines the probability of an event $\mathscr{E}$ occurring at a specific flight k, given the state of the system x, characterized by the evolution of the CF. It is now possible to calculate the trend of the variable τ_ε, which indicates the flight after which the event $\mathscr{E}$ is most probable.

As carried out in [122], $x^{(i)}_{k_p+1:k} = \left\{x^{(i)}_j\right\}^k_{j=k_p+1}$ indicates the i-th trace from flight k_p to cycle k. i indicates a specific trace and, as a consequence, $i \in 1, 2, ..., N_P$, since there are N_P particles and traces.

The probability that the first event $\mathscr{E}$ occurs at flight k can be denoted as $\mathbb{P}(\tau_\varepsilon = k)$ and calculated according to Equation 5.17 from [122].

$$\mathbb{P}(\tau_\varepsilon = k) = \frac{1}{N} \sum_{i=1}^{N} \left[\mathbb{P}\left(E_k = \mathscr{E} \mid x^{(i)}_k\right) \cdot \prod_{j=k_p+1}^{k-1} \left(1 - \mathbb{P}\left(E_k = \mathscr{E} \mid x^{(i)}_j\right)\right) \right] \tag{5.17}$$

The PoF is hence calculated using the mathematical formulation reported in Equation 5.17. Once the PoF trend is obtained, the risk trend can be computed by the integration of the PoF trend in time, resulting in a curve which has a 0 to 100 range. From both a mathematical and management standpoint, incorporating risk values into decision-making processes is considered as the most effective approach for managing uncertainty. Risk quantifies the likelihood of an UR occurring, and when risk reaches a value of 100 percent, it signifies absolute certainty of that event having occurred. The user has the option to select their desired risk level; in the rest of the paper the risk target value is set as 99 per cent. The Acuña risk trend [122] is then passed to the Health Management, Strategies & Support Layer.

An example is provided in Figure 5.31, where the results of the Probability of Failure and Risk Computation Module are reported. The same CF "Work_RHT_Skew" is considered and plotted against the PAS flight number. The PAS, SN 47, had already suffered one UR at flight 291 and a prediction is performed a few flights after the first UR ($k_p = 296$). The results of the MC simulations are reported and highlight an excellent simulation ability, noticeable by comparing the distributions in the blue "fanchart" with the GT after k_p. The second UR happened at flight 451, as highlighted by the vertical thick orange line. The green trend represents the PoF

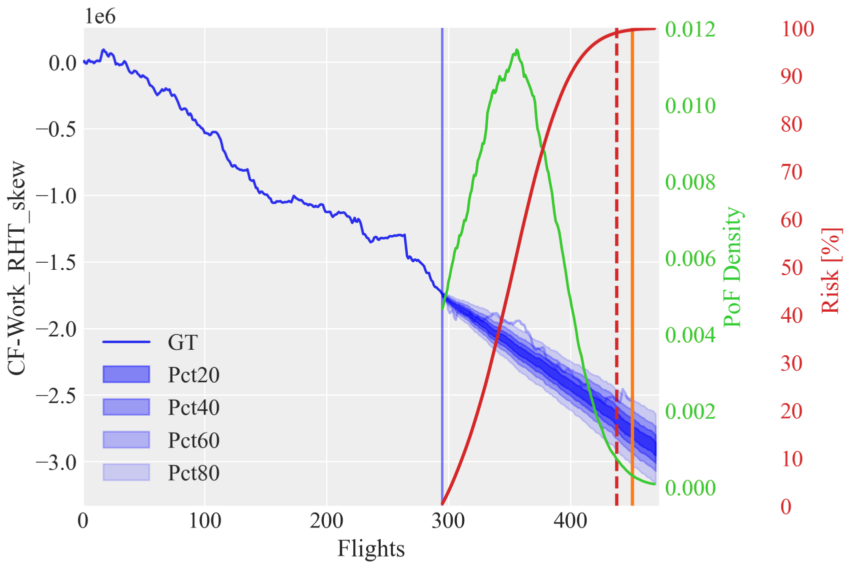

Fig. 5.31 MC simulations and PoF/Risk results with 1000 particles for as many as 200 simulated flights starting at flight $k_p = 297$. The GT after k_p is reported and it is inside the prediction "cone", marked by the 80 percentile. The PAS SN already suffered one UR at 291 flights and the second UR occurred at flight 451, as shown by the vertical orange line. A sharp increase in risk can be noted due to the fact that, at k_p, the PoF is already near its peak. The prediction associated with 99 per cent risk is highlighted by the vertical dotted red line.

density, while the red curve trend represents the risk. It should be noted that the risk reaches a high percentage before the actual UR occurred which implies that, although precision can be improved with more data, failure is not missed (or underestimated), thus raising confidence in the CF processing. A red dotted line is reported when risk reaches 99 per cent, indicating a good estimate of the possible next UR. Further details involving the evaluation through a literature established metric are reported in Section 5.4.1.

5.4.1 Metrics & Discussion

A typical challenge related to the development and acceptance of PHM results is linked to the wide range of applicable validation metrics. To make the analysis more grounded and demonstrate the framework solidity and future potential, the author decided to split the evaluation phase into two parts: the first linked to the Probability Health Prediction Module to evaluate how well the predictions compare to the GT and the second part which assesses the PoF and Risk Computation Module performance.

Probabilistic Health Prediction Module Performance Evaluation

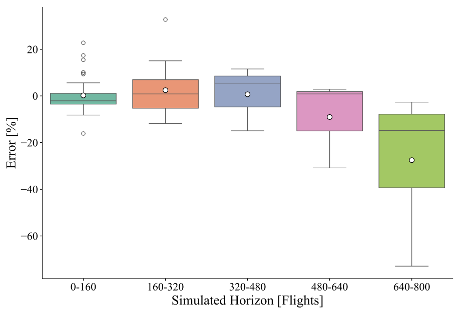

(a) Mean relative error between the CFs future trends simulations and the GT.

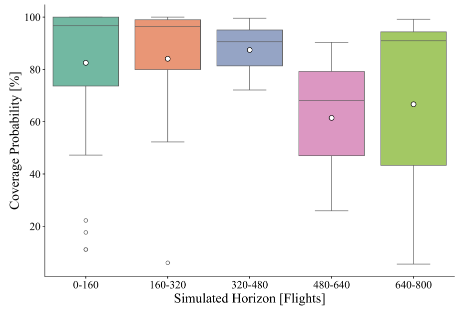

(b) Coverage probability inside the 95-th percentile.

Fig. 5.32 Box plot results of the metrics for the Probabilistic Health Predictions Module evaluated at different simulation horizons across the whole PAS fleet. The median is represented with the horizontal line inside the box, whereas the white dot inside the box represents the mean. Both the mean relative error and the coverage probability highlight the capability of this module in predicting the CF trend, with better performance in the medium term.

Figures 5.32a and 5.32b show the statistics of the evaluation metrics chosen for the Probabilistic Health Prediction Module, performed across the whole PAS fleet at different simulated horizons. In particular, Figure 5.32a reports the box plot of the mean relative error calculated between the mean of the probability cone and the GT and shows a prediction performance with limited relative errors within acceptability margins defined by the OEM, especially in the mid term. When the simulations are carried out for more than 640 flights the mean of the error, despite being higher, still is lower than 20 per cent. Figure 5.32b shows the probability that the 95th percentile cone captures the actual GT. The results show extremely promising performance in the short and mid horizon. In the 0-160 flight band the presence of some outliers can be noted. This is because the prediction horizon is short and the probability that the GT is outside of the cone is higher because the cone is very narrow. However, the error in the same horizon band confirms that the GT does not differ much from the prediction. These metrics confirm the Probabilistic Health Predictions Module ability to simulate Cumulative Features (CFs) behavior.

PoF and Risk Computation Module Performance Evaluation

The next paragraphs report the evaluation metrics for the Probability of Failure and Risk Computation Module. The validation metric is calculated using a validation subset of URs that meet two criteria: they must have occurred in the latter part of the recorded data-set, and their reported causes must be directly related to prognosable failures (in accordance with the information provided). These criteria exclude older records, where data storing was not performed systematically and failures due to unprognosable/unkown events (e.g., no information reported, vibration, actuator locked, external leakage, etc.). Moreover, the predictions were performed by excluding the final CF value associated with each UR from its respective threshold distribution. This method is similar to cross-validation, preventing contamination between training and testing phases.

Multiple studies have tried to standardize the metrics panorama and a possible solution is represented by the widely recognized $\alpha - \lambda$ curves, as explained in Saxena et al. in [146, 147]. $\alpha - \lambda$ curves compare the actual RUL against the predicted RUL. α refers to the acceptance bound of the uncertainty cone (typically set at 20 %), while λ indicates the normalized time scale defined as $\lambda = \frac{k}{k_f}$ where k is the present time and k_f indicates the time-to-failure. λ is hence plotted on the x-axis. When λ equals one, the equipment has reached the UR. A normalized version is often selected to compare different traces in a single graph: in this case various SN with different age. Figure 5.33a shows an example of $\alpha - \lambda$ for a specific SN, whereas Figure 5.33b shows the $\alpha - \lambda$ plot obtained for a set of 12 predictions in the validation subset. The mean value is reported in blue along with the uncertainty indication in a box-plot fashion at specific values of λ.

Despite the wide threshold PDF variance, Figure 12b shows promising results as the mean is always inside the prediction cone. The system does converge for the RUL estimation and, although suboptimal in terms of precision, provides useful information from a maintenance perspective. This is consistent with the primary objective of the framework: offering a reasonable time horizon for maintenance personnel, rather than seeking the exact point of failure to perform just-in-time maintenance. Moreover, it has to be noted that in most of the cases, the framework is underestimating the actual RUL, showing to be conservative. This underlines once again that the methodology is robust and is able to track degradation patters: it is hence deemed that, once new data is logged and the threshold PDF is more precise,

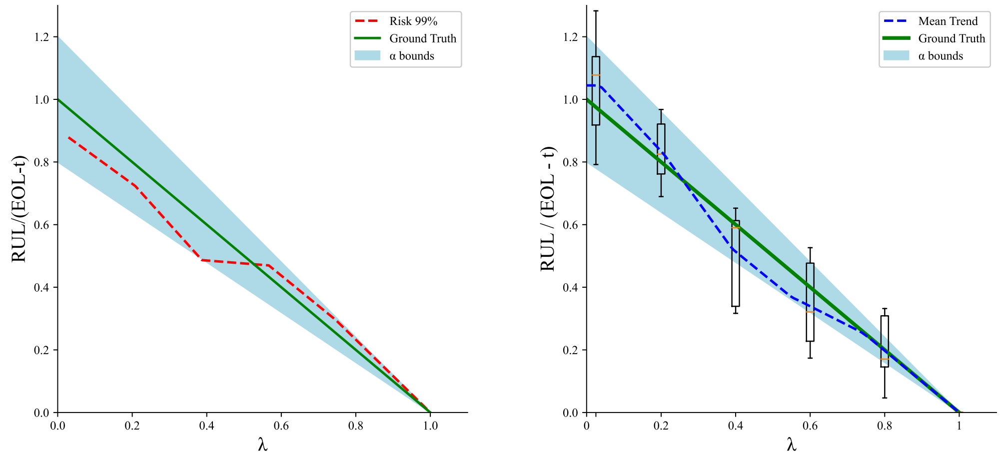

(a) $\alpha - \lambda$ plot for PAS 57 when risk reaches 99 per cent.

(b) $\alpha - \lambda$ plot across the subset of prognosable URs with uncertainty quantification in a box-plot fashion.

Fig. 5.33 Metric for the Probability of Failure and Risk Computation Module: $\alpha - \lambda$ plots calculated when risk reaches 99 per cent.

the methodology will provide even better results, thus reducing the variance and the linked inter quartile range especially at λ of 0.8 and higher. Since the degradation processes are often specific to the particular failure modes, an additional step was undertaken to reduce volatility in the hazard threshold. Each UR was classified according to its specific removal reasons as already reported in Figure 4.8, and the top five CF histograms were reviewed to identify potential patterns or sub-histograms within the overall threshold.

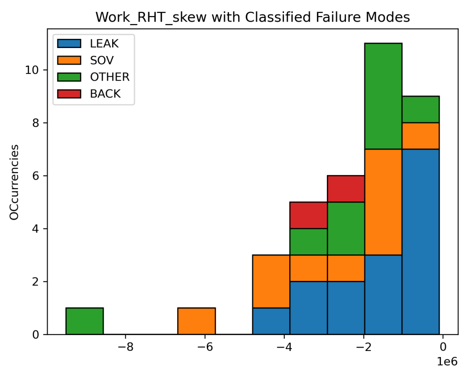

Fig. 5.34 Fault Mode analysis on the "Work RHT skew" CF.

For example, Figure 5.34 illustrates this detailed analysis for the RHT mechanical work threshold using kurtosis. The histogram presented here corresponds to the one shown in Figure 5.4, with colors indicating different failure modes. Unfortunately, no clear patterns were identified, suggesting that further reduction of variance in the resulting probability density function may not be feasible at this stage.

Chapter 6

Layer 3: Health Management, Strategies, and Support

As already discussed in Section 5.3, the design of the Probabilistic Health Prediction Module is inherently modular, enabling seamless interaction with the Health Management, Strategies & Support Layer. This architecture allows both short- to mid-term and long-term optimization strategies to be derived from a single, coherent risk evaluation pipeline, thus supporting tailored maintenance and operational solutions that can be adapted to different customer requirements and fleet contexts. At the same time, it must be stressed that the effective implementation of any optimization strategy in the field depends critically on the specific maintenance policies, regulatory constraints, contractual agreements, and actual operational use of the assets. For this reason, the present section does not prescribe a single practical implementation path, nor does it claim to define a universal "optimal" policy. Instead, it discusses a set of possible integration scenarios, illustrating how the outputs of the proposed framework could be embedded within real-world maintenance planning and fleet management processes, while leaving the detailed configuration and deployment to the specific operational context in which the asset is employed.

6.1 Short to Mid-Term Methodologies

In the short- to mid-term scenario, the module is used as a tactical decision-support tool by simulating a set of predefined CMTs and assessing the resulting evolution of the risk of URs. It should be noted that there is no direct forward correspondence between a CMT and a unique mission type, as discussed previously, since different missions may generate similar flight profiles and the clustering was performed on signal trends rather than on mission labels. Nevertheless, if fleet managers or customer training staff can provide an expected flight profile for an upcoming mission, even in approximate form, this profile can be processed by the trained SOM and associated with the most similar neuron, and thus with the corresponding CMT. In this sense, the SOM does not identify the nominal mission category itself, but rather maps the expected operational profile onto the closest cluster of historically observed flight behaviors. More generally, this also opens the possibility of adopting a soft assignment logic, in which a new profile is associated not only with a single CMT, but with multiple CMTs weighted according to their degree of similarity (e.g., with Gaussian Mixture Models (GMMs)), thereby supporting a more flexible representation. This "what-if" capability allows the operator to evaluate how alternative mission sequences or tail assignments would impact the health and risk profile of the monitored actuator.

On this basis, several practical decision logics can be built. For example, the system can provide tail-number recommendations by ranking the available aircraft for a given mission according to the incremental risk they incur, favoring tails with healthier actuators for demanding mission profiles, and reserving more degraded assets for less critical or shorter flights. Similarly, health-informed metrics can support repositioning decisions (e.g., moving a tail with an at-risk component towards an operating base with better maintenance capability or spare availability) or advise opportunistic removals of PAS units when an upcoming sequence of severe missions is foreseen and the predicted risk crosses operator-defined thresholds. These ideas are consistent with recent work on integrating prognostics and CBM into fleet maintenance scheduling, tail assignment, and airline decision-support frameworks, where health information is embedded into optimization models for maintenance scheduling and fleet assignment. Likewise, PHM-driven risk indicators have been shown to be valuable for spare parts planning and condition-based ordering policies, where predicted RUL and associated uncertainty guide when and where to position critical components.

The key point is that all these tactical uses emerge without introducing new sensors, by exploiting aircraft-level data to derive virtual health indicators for the PAS and feeding them into existing operational research and planning processes.

The proposed framework therefore provides the missing quantitative link between mission planning, tail assignment, and component-level risk, while remaining modular enough to interface with different airline policies, network structures, and optimization tools.

6.2 Long-Term Methodologies

In a longer-term perspective, the same framework can be operated in an "open-loop" fashion by feeding it with prospective CMT sequences whose statistics are derived from historical mission data. In practice, long-horizon usage scenarios can be generated by sampling CMT sequences from empirically estimated distributions (e.g., per squadron, training phase, season, or theatre of operation) and propagating the corresponding CFs and hazard PDFs over an extended time horizon. This allows the operator to evaluate how different long-term usage profiles (more aggressive training syllabi, changes in deployment patterns, introduction of new mission types) impact the aggregate risk of URs and removals at fleet level. The resulting risk trends can then be exploited to support strategic maintenance optimization, spare parts planning, and storage/repair capacity sizing, by providing quantitative forecasts of expected failures and associated uncertainty over months or years rather than single rotations.

A natural extension is the integration of this risk information within broader logistics and lifecycle cost (LCC) frameworks. For example, predicted distributions of future removals can be coupled with models of spare provisioning and inventory control (min–max policies, base- and pool-level stocks, repair turnaround times) to dimension the number and location of PAS spares needed to achieve desired service levels and dispatch reliability targets. Likewise, long-term health predictions can inform capacity planning for maintenance technicians, hangar slots, and test benches, aligning manpower and infrastructure with the expected workload rather than relying solely on historical averages or fixed-interval inspections. In a contractual setting, such as power-by-the-hour or availability-based contracts, the same risk outputs can be shared with the OEM to support joint decisions on overhaul strategies, repair vs replace policies, and design improvements, thus linking PHM with LCC and reliability-centered maintenance at a more strategic level.

The specific actions implemented within the Health Management, Strategies & Support Layer must, however, be designed around how the AJT operator actually manages and operates the fleet. The exploitation of long-term risk trends should be consistent with existing planning horizons (e.g., seasonal training plans, annual budget cycles) and with the way missions and syllabus are structured. A practical and intuitive implementation is to establish a direct mapping between CMTs and the syllabus missions flown by trainee pilots: in this way, fleet managers can simulate the impact of alternative training plans (e.g., front-loading demanding manoeuvres vs distributing them over time) on actuator health and associated future removals. Similar logic can be extended to deployment strategies, base assignments, and rotation policies, enabling scenario analyses that combine operational effectiveness with equipment preservation.

Chapter 7

Conclusions

Developing PHM systems capable of detecting abnormal behaviors and enhancing maintenance management throughout the asset lifecycle is a complex yet rewarding endeavor. This thesis has addressed the problem of implementing a PHM system for a legacy AJT HT EHA, using only aircraft-level operational data and without any modification to the existing hardware or avionics architecture. The work was carried out under constraints that are typical of in-service aerospace platforms: limited and low-frequency measurements, heterogeneous historical records, and strict certification implications for any design change. Within this context, the thesis has demonstrated that it is possible to extract meaningful health information from non-dedicated data streams and to potentially convert it into decision support for maintenance planning.

The core contribution of this thesis lies in the design, implementation, and evaluation of a PHM architecture that operates entirely on in-service data, sampled at low frequency and collected for purposes other than PHM. The proposed framework adopts a layered structure of functional modules, starting from raw flight and maintenance data and progressing through feature extraction, health indicator construction, probabilistic projection, and risk-oriented interpretation. Using statistical descriptors of cumulative loads and usage, combined with Monte Carlo simulations and clustering techniques, the system produces probabilistic estimates of failure timeframes and associated risk trends for the EHA fleet. Importantly, this is achieved without relying on local condition-monitoring signals and without modifying aircraft hardware or existing data management pipelines.

From a scientific and technical standpoint, the thesis offers several contributions. First, it provides a documented, end-to-end PHM framework tailored to legacy aerospace platforms, bridging the gap between theoretical PHM concepts and the realities of industrial implementation. Second, it shows that meaningful prognostic information can be derived from aggregate aircraft-level data, provided that the underlying physics of use and the operational context are carefully considered in the design of health indicators. Third, it demonstrates that probabilistic prognostics, via Monte Carlo simulation and clustering of projected behaviours, can be integrated into a practical decision-support tool, yielding conservative yet actionable estimates of remaining life and risk evolution. Finally, the work openly discusses and quantifies the impact of incomplete failure information, limited data coverage, and legacy constraints, thereby contributing a realistic benchmark for future PHM developments in similar environments.

The metrics used to evaluate the results highlight both the strengths and the current limits of the proposed approach. The predictions are deliberately conservative, reflecting the scarcity of confirmed failure events and the incomplete description of degradation mechanisms. Uncertainty bands remain relatively wide, especially in the tails of the distributions. Yet, even under these conditions, the framework is capable of distinguishing higher-risk units from those with safer usage profiles, of identifying emerging trends at fleet level, and of providing a quantitative basis for maintenance planning decisions. In a real industrial context, where decisions are often made under uncertainty and with imperfect information, such conservative prognostic outputs can still offer practival value by supporting prioritization, resource allocation, and early warning.

A key aspect of this work is that all design choices were consistently aligned with the constraints of a legacy platform: long service life, limited sensorization, fixed avionics configuration, strict certification implications of any modification, and fragmented historical records. Rather than treating these constraints as obstacles that preclude PHM, the thesis has treated them as design requirements. The resulting framework is robust, relatively simple to integrate, and conceptually transparent, features that are crucial for acceptance by operators, maintainers, and regulators. In this regard, the thesis not only presents a novel algorithm but also provides a practical framework for integrating PHM into an existing ecosystem, addressing a common and often overlooked challenge within the aerospace industry.

The long-term impact of this approach extends beyond the specific AJT case. The methodology is naturally compatible with other military and civil fleets where mission profiles are well characterized and operational data are routinely collected, even if at modest sampling rates. In civil aviation, where data volumes, fleet sizes, and operational regularity are typically higher, the same principles could be applied with even greater effectiveness. As more flight hours are flown and more events (both failures and near-misses) are recorded, the statistical power of the framework will increase. Over the aircraft life cycle, this will allow progressive refinement of health indicators, calibration of probabilistic models, and, ultimately, a more precise quantification of risk and remaining useful life. In other words, the framework is designed not as a static tool, but as a growing system that learns from the evolving experience of the fleet.

Moreover, the thesis opens several lines for future research and industrial development. On the methodological side, the statistical features and Monte Carlo projections developed here can be combined with more advanced machine learning or hybrid physics-based/data-driven or even physics-informed models [148, 149], provided that interpretability and robustness are preserved. Uncertainty quantification can be further refined through more sophisticated probabilistic metrics and calibration techniques, enabling a more rigorous assessment of prognostic quality and confidence. On the system level, integrating the outputs of this PHM framework with maintenance planning tools, logistic support analyses, and cost models would allow moving from purely technical prognostics to fully fledged predictive decision-making, where technical risk and economic impact are evaluated jointly.

From an industrial perspective, the thesis suggests a pragmatic roadmap for organizations that wish to introduce PHM capabilities without embarking on costly retrofit programmes. By showing that useful health information can be extracted from existing data streams, the work lowers the perceived barrier to entry for PHM and encourages incremental deployment. Fleet operators can start with limited-scope applications, such as risk ranking of units, monitoring of critical sub-systems, or support for campaign-based maintenance decisions, and progressively expand the scope as data and experience accumulate. This incremental approach is particularly relevant in industries where certification cycles are long, budgets are constrained, and the tolerance for operational disruption is low.

Finally, at a broader level, the thesis contributes to the ongoing evolution of PHM from a promising research area to a mature component of asset management in aerospace and related sectors. It does so by embracing the complexity and imperfections of real industrial data, by explicitly acknowledging the compromises inherent in legacy environments, and by demonstrating that meaningful gains can still be achieved under these conditions. The work shows that PHM need not be reserved only for newly designed “smart” platforms; it can also be retrofitted, conceptually, onto existing fleets through careful exploitation of the data that are already available.

In conclusion, this thesis has provided a tested and documented PHM framework for legacy aerospace systems and, equally importantly, a transparent discussion of the design decisions, limitations, and trade-offs involved in its development. It offers practical guidance for implementing prognostics in complex environments characterized by limited and imperfect data; conditions that are, in fact, the norm rather than the exception in aeronautics and many other industries. While the present results are necessarily constrained by the available information, they delineate a credible and validated pathway towards data-driven predictive maintenance for long-lived platforms. As such, the work lays a foundation on which future research and industrial initiatives can build, progressively advancing PHM from isolated case studies to a standard, trusted component of aerospace maintenance practice.

References

[1] Leonardo Baldo, Andrea De Martin, Mathieu Terner, Giovanni Jacazio, and Massimo Sorli. Implementing phm for legacy flight control actuators through operational aircraft data: Approach and lessons learned. *Results in Engineering*, 28:107214, 2025.

[2] Francesco Pace. Sviluppo di modelli della dinamica di attuatori per comandi primari di velivoli fly-by-wire. Tesi di laurea specialistica, Università di Pisa, Pisa, Italy, April 2008. Corso di laurea: Ingegneria Aerospaziale. Relatori: Prof. Eugenio Denti; Ing. Mauro Norese; Ing. Germana Vinelli; Prof. Roberto Galatolo; Dott. Gianpietro Di Rito.

[3] Leonardo Baldo, Andrea De Martin, Giovanni Jacazio, and Massimo Sorli. A systematic literature review on phm strategies for (hydraulic) primary flight control actuation systems. *Actuators*, 14(8), 2025.

[4] Leonardo Baldo, Andrea De Martin, Massimo Sorli, and Mathieu Terner. Condition-based-maintenance for fleet management. In *Aerospace Science and Engineering - III Aerospace PhD-Days*, pages 57–60, 2023.

[5] Leonardo Baldo, Andrea De Martin, Mathieu Terner, and Massimo Sorli. The journey towards condition-based maintenance: A framework for the horizontal tail actuator of an advanced jet trainer aircraft. In *34th Congress of the International Council of the Aeronautical Sciences (ICAS)*, 2024.

[6] Leonardo Baldo. Development of a data-driven condition-based maintenance methodology framework for an advanced jet trainer. In *PHM Society European Conference*, volume 8, pages 5–5, October 2024.

[7] Leonardo Baldo, Andrea De Martin, Mathieu Terner, Giovanni Jacazio, Marcos E Orchard, and Massimo Sorli. Mission profile clustering for usage-based health modeling of flight control actuators applied to a fleet of advanced jet trainers. In *Annual Conference of the PHM Society*, volume 17, 2025. Best Paper Award.

[8] Harry A Kinnison and Tariq Siddiqui. *Aviation maintenance management*. McGraw-Hill Education, 2013.

[9] European Commission. Commission regulation (eu) no 1321/2014 of 26 november 2014 on the continuing airworthiness of aircraft and aeronautical products, parts and appliances, and on the approval of organisations and personnel involved in these tasks text with eea relevance. Technical report, European Commission, 2014.

[10] Enrico Zio. Prognostics and health management (phm): Where are we and where do we (need to) go in theory and practice. *Reliability Engineering & System Safety*, 218:108119, 2022.

[11] Gartner. Market trends: Predictive maintenance drives iot in manufacturing operations. https://www.gartner.com/en/documents/3856379, 2018. Accessed: 5th March 2025.

[12] ICF International Inc. Using digital threads and twins to improve the mro lifecycle. https://www.icf.com/insights/transportation/improving-mro-lifecycle, 2020. Accessed: 5th March 2025.

[13] Leonardo S.p.A. Piano industriale - aggiornamento 2025 (2025–2029). https://www.leonardo.com/it/investors/industrial-plan, 2025. Accessed: 20th January 2026.

[14] Leonardo S.p.A. Le tecnologie multi-dominio per affrontare gli scenari operativi del futuro. https://www.leonardo.com/it/focus-detail/-/detail/tecnologie-multi-dominio_focus, July 2024. Accessed: 20th January 2026.

[15] Hamad Al-kaabi, Andrew Potter, and Mohamed Naim. An outsourcing decision model for airlines' mro activities. *Journal of Quality in Maintenance Engineering*, 13(3):217–227, 2007.

[16] G. Smith, J.B. Schroeder, S. Navarro, and D. Haldeman. Development of a prognostics and health management capability for the Joint Strike Fighter. In *1997 IEEE Autotestcon Proceedings AUTOTESTCON '97. IEEE Systems Readiness Technology Conference*, pages 676–682, 1997.

[17] A. Hess and L. Fila. The Joint Strike Fighter (JSF) PHM concept: Potential impact on aging aircraft problems. In *Proceedings, IEEE Aerospace Conference*, volume 6, pages 6–6, 2002.

[18] Andrew K. S. Jardine, Daming Lin, and Dragan Banjevic. A review on machinery diagnostics and prognostics implementing condition-based maintenance. *Mechanical Systems and Signal Processing*, 20(7):1483–1510, 2006.

[19] Luis Basora, Paloma Bry, Xavier Olive, and Floris Freeman. Aircraft fleet health monitoring with anomaly detection techniques. *Aerospace*, 8(4):103, 2021.

[20] Mohamed Cherif Dani, Cassiano Freixo, François-Xavier Jollois, and Mohamed Nadif. Unsupervised anomaly detection for aircraft condition monitoring system. In *2015 IEEE Aerospace Conference*, pages 1–7, 2015.

[21] H.K. Ng, R.H. Chen, and J.L. Speyer. A vehicle health monitoring system evaluated experimentally on a passenger vehicle. *IEEE Transactions on Control Systems Technology*, 14(5):854–870, 2006.

[22] Andrews Darfour Kwakye, Ian K Jennions, and Cordelia Mattuvarkuzhali Ezhilarasu. Platform health management for aircraft maintenance – a review. *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, 238(3):267–283, 2024.

[23] Andy J Evans. Managing a successful hums operation. In *Third International Conference on Health and Usage Monitoring-HUMS2003*, page 109. Citeseer, 2002.

[24] J.E. Land. Hums-the benefits-past, present and future. In *2001 IEEE Aerospace Conference Proceedings (Cat. No.01TH8542)*, volume 6, pages 3083–3094 vol.6, 2001.

[25] Joanna Kappas and Peter Frith. From hums to phm: Are we there yet. In *Proceedings of the 17th Australian International Aerospace Congress, Melbourne, VIC, Australia*, pages 26–28, 2017.

[26] Ricardo de Arriba and Alberto Gallego. Health and usage monitoring system (hums) strategy to enhance the maintainability & flight safety in a flight control electromechanical actuator (ema). In *PHM Society European Conference*, volume 3, 2016.

[27] Spherical Insights. Global aerospace defense mro market insights forecasts. https://www.sphericalinsights.com/reports/aerospace-defense-mro-market, 2023. Accessed: 5th March 2025.

[28] Mordor Intelligence. Military aviation mro market size & share analysis - growth trends & forecasts (2025 - 2030). https://www.mordorintelligence.com/industry-reports/military-aviation-maintenance-repair-and-overhaul-market, 2024. Accessed: 5th March 2025.

[29] Mordor Intelligence. Military aviation market size & share analysis - growth trends & forecasts up to 2030. https://www.mordorintelligence.com/industry-reports/military-aircraft-market, 2024. Accessed: 5th March 2025.

[30] Lester O Patterson and Bardford C Tonder. *External strategic analysis of the aviation maintenance, repair and overhaul (MRO) industry and potential market opportunities for Fleet Readiness Center Southwest*. PhD thesis, Naval Postgraduate School Monterey, CA, USA, 2009.

[31] Defense Systems Management College. *Acquisition Logistics Guide*. US Government Printing Office, 1997.

[32] Office of Cost Assessment and Program Evaluation. Operating and support cost-estimating guide. Technical report, U.S. Department of Defense, January 2025.

[33] Scott C Hewitson, Jonathan D Ritschel, Edward White, and Gregory Brown. Analyzing operating and support costs for air force aircraft. *Journal of Defense Analytics and Logistics*, 2(1):38–54, 2018.

[34] Garrett B O'Hanlon, Jonathan D Ritschel, Edward D White, and Gregory E Brown. Delineating operating and support costs in aircraft platforms. *Defense Acquisition Research Journal*, 2018.

[35] Gary L Jones, Edward D White, Erin T Ryan, and Jonathan D Ritschel. Investigation into the ratio of operating and support costs to life-cycle costs for dod weapon systems. *Defense Acquisition Research Journal*, 2014.

[36] AW&ST. 2013. The mro global leading players. Technical report, Aviation Week and Space Technology, MRO Edition, Aviation Week Ranks, 2013.

[37] Amy Mainville Cohn and Cynthia Barnhart. Improving crew scheduling by incorporating key maintenance routing decisions. *Operations Research*, 51(3):387–396, 2003.

[38] Le pôle interministériel de prospective et d'anticipation des mutations économiques (PIPAME). Maintenance et réparation aéronautique. Technical report, PIPAME, 2010.

[39] Basem Yousef Badkook. Study to determine the aircraft on ground (aog) cost of the boeing 777 fleet at x airline. *American Scientific Research Journal for Engineering, Technology, and Sciences*, 25:51–71, 2016.

[40] K. B. Öner, R. Franssen, G. P. Kiesmuller, and G. J. van Houtum. Life cycle costs measurement of complex systems manufactured by an engineer-to-order company. Technical report, Technische Universiteit Eindhoven, 2007.

[41] Wim JC Verhagen, Lennaert WM De Boer, and Richard Curran. Component-based data-driven predictive maintenance to reduce unscheduled maintenance events. In *Transdisciplinary Engineering: A Paradigm Shift*, pages 3–10. IOS Press, 2017.

[42] Nima Safaei, Dragan Banjevic, and Andrew KS Jardine. Workforce-constrained maintenance scheduling for military aircraft fleet: a case study. *Annals of Operations Research*, 186:295–316, 2011.

[43] Gregory W. Vogl, Brian A. Weiss, and M. Alkan Donmez. Standards Related to Prognostics and Health Management (PHM) for Manufacturing. In *Annual Conference of the PHM Society*, volume 6, 2014.

[44] ISO 17359:2018 condition monitoring and diagnostics of machines — general guidelines. Standard, International Organization for Standardization, 2018.

[45] Avin Mathew, Liqun Zhang, Sheng Zhang, and Lin Ma. A review of the mimosa osa-eai database for condition monitoring systems. In Joseph Mathew, Jim Kennedy, Lin Ma, Andy Tan, and Deryk Anderson, editors, *Engineering Asset Management*, pages 837–846, London, 2006. Springer London.

[46] ISO 13374-1:2003 Condition monitoring and diagnostics of machines — Data processing, communication and presentation Part 1: General guidelines. Standard, International Organization for Standardization, 2003.

[47] ADS-79E-HDBK Condition Based Maintenance System for US Army Aircraft. Handbook, US Army, 2016.

[48] OSA-CBM Open System Architecture for Condition-Based Maintenance. Technical report, MIMOSA, 2016.

[49] SAE ARP 6883:2019 Guidelines for Writing IVHM Requirements for Aerospace Systems. Aerospace recommended practice, Society of Automotive Engineers, 2019.

[50] Abhinav Saxena, Kai Goebel, Don Simon, and Neil Eklund. Damage propagation modeling for aircraft engine run-to-failure simulation. In *2008 International Conference on Prognostics and Health Management*, pages 1–9, October 2008.

[51] Bhaskar Saha, Kai Goebel, Scott Poll, and Jon Christophersen. Prognostics methods for battery health monitoring using a Bayesian framework. *IEEE Transactions on Instrumentation and Measurement*, 58(2):291–296, February 2009.

[52] Jokin Alcibar, Jose I Aizpurua, Ekhi Zugasti, Carmen Alonso-Montes, and Ibon Diez. Towards a probabilistic error correction approach for improved drone battery health assessment. In *e-proceedings of the 33rd European Safety and Reliability Conference (ESREL 2023), Southamp-ton, United Kingdom, on 3–7 September 2023*, 2023.

[53] Jokin Alcibar, Jose I Aizpurua, and Ekhi Zugasti. Towards a probabilistic fusion approach for robust battery prognostics. In *PHM Society European Conference*, volume 8, pages 13–13, 2024.

[54] Bruno Masserano, Jorge E. Garcia Bustos, Camilo Ramirez, Benjamin Brito Schiele, Cristobal E. Allendes, Ricardo Salas-Espineira, Sofia Mancilla, Jose Luis Espinoza, Aramis Perez, Francisco Jaramillo-Montoya, and Marcos E. Orchard. A lithium-ion battery degradation model agnostic to cell chemistry with integrated state-of-charge and temperature dependence. In *Annual Conference of the PHM Society*, volume 17, 2025.

[55] José Joaquín Mendoza Lopetegui, Gianluca Papa, and Mara Tanelli. Active and data-driven health and usage monitoring of aircraft brakes. *IFAC-PapersOnLine*, 56(2):10740–10745, 2023.

[56] Shafayat Hasan Chowdhury, Fakhre Ali, and Ian K. Jennions. A methodology for the experimental validation of an aircraft ECS digital twin targeting system level diagnostics. In *Annual Conference of the PHM Society*, volume 11, 2019.

[57] Jokin Alcibar, Jose I. Aizpurua, Ekhi Zugasti, and Oier Peñagarikano. A hybrid probabilistic battery health management approach for robust inspection drone operations. *Engineering Applications of Artificial Intelligence*, 146:110246, 2025.

[58] Bulent Sarlioglu and Casey T. Morris. More Electric Aircraft: Review, Challenges, and Opportunities for Commercial Transport Aircraft. *IEEE Transactions on Transportation Electrification*, 1(1):54–64, June 2015.

[59] J. A. Rosero, J. A. Ortega, E. Aldabas, and L. Romeral. Moving towards a more electric aircraft. *IEEE Aerospace and Electronic Systems Magazine*, 22(3):3–9, 2007.

[60] Wen-Bin Chen, Xiao-Yang Li, and Rui Kang. Attractor based performance characterization and reliability evolution for electromechanical systems. *Mechanical Systems and Signal Processing*, 222:111803, 2025.

[61] Zhengyang Yin, Niaoqing Hu, Jiageng Chen, Yi Yang, and Guoji Shen. A review of fault diagnosis, prognosis and health management for aircraft electromechanical actuators. *IET Electric Power Applications*, 16(11):1249–1272, 2022.

[62] Seyed Mohammad Rezvanizaniani, Jacob Dempsey, and Jay Lee. An Effective Predictive Maintenance Approach based on Historical Maintenance Data using a Probabilistic Risk Assessment: PHM14 Data Challenge. *International Journal of Prognostics and Health Management*, 5(2), 2014.

[63] Hyunjae Kim, Taewan Hwang, Jungho Park, Hyunseok Oh, and Byeng D Youn. Risk prediction of engineering assets: An ensemble of part lifespan calculation and usage classification methods. *International Journal of Prognostics and Health Management*, 5(2), 2014.

[64] Luuk Schoenmakers. *Condition-based Maintenance for the RNLAF C-130H(-30) Hercules*. Master's thesis, Eindhoven University of Technology, November 2020.

[65] Juncheng Mi and Guoqin Huang. Dynamic Prediction of Performance Degradation Characteristics of Direct-Drive Electro-Hydraulic Servo Valves. *Applied Sciences*, 13(12):7231, January 2023.

[66] Carl S. Byington, Matthew Watson, and Doug Edwards. Data-driven neural network methodology to remaining life predictions for aircraft actuator components. In *2004 IEEE Aerospace Conference Proceedings (IEEE Cat. No.04TH8720)*, volume 6, pages 3581–3589, March 2004.

[67] Vignesh V. Shanbhag, Thomas J. J. Meyer, Leo W. Caspers, and Rune Schlanbusch. Failure monitoring and predictive maintenance of hydraulic cylinder—state-of-the-art review. *IEEE/ASME Transactions on Mechatronics*, 26(6):3087–3103, 2021.

[68] Antonio C. Bertolino, Rocco Gentile, Giovanni Jacazio, Francesco Marino, and Massimo Sorli. EHSA Primary Flight Controls Seals Wear Degradation Model. In *ASME International Mechanical Engineering Congress and Exposition*, volume Volume 1: Advances in Aerospace Technology, page V001T03A024, January 2018.

[69] Qun Chao, Yuechen Shao, Chengliang Liu, and Xiaoxue Yang. Health evaluation of axial piston pumps based on density weighted support vector data description. *Reliability Engineering & System Safety*, 237:109354, 2023.

[70] Alberto Bacci, Antonio C. Bertolino, Andrea De Martin, and Massimo Sorli. Multiphysics Modeling of a Faulty Rod-End and Its Interaction With a Flight Control Actuator to Support PHM Activities. In *ASME 2021 International Mechanical Engineering Congress and Exposition*. American Society of Mechanical Engineers Digital Collection, 2022.

[71] Wlamir OL Vianna and João Pedro P Malere. Aircraft hydraulic system leakage detection and servicing recommendations method. In *Annual Conference of the PHM Society*, volume 6, 2014.

[72] Wenhao Zong, Fangyi Wan, and Yishui Wei. Real-time monitoring for the actuator mechanism of the aileron. In *2017 Prognostics and System Health Management Conference (PHM-Harbin)*, pages 1–5, 2017.

[73] Hongmei Liu, Jichang Zhang, and Chen Lu. Performance degradation prediction for a hydraulic servo system based on elman network observer and gmm–svr. *Applied Mathematical Modelling*, 39(19):5882–5895, 2015.

[74] *Prognostics and Health Monitoring of Electro-Hydraulic Systems*, volume Volume 2: Mechatronics; Estimation and Identification; Uncertain Systems and Robustness; Path Planning and Motion Control; Tracking Control Systems; Multi-Agent and Networked Systems; Manufacturing; Intelligent Transportation and Vehicles; Sensors and Actuators; Diagnostics and Detection; Unmanned, Ground and Surface Robotics; Motion and Vibration Control Applications of *Dynamic Systems and Control Conference*. American Society of Mechanical Engineers Digital Collection, November 2017.

[75] Chen Lu, Hang Yuan, and Jian Ma. Fault detection, diagnosis, and performance assessment scheme for multiple redundancy aileron actuator. *Mechanical Systems and Signal Processing*, 113:199–221, December 2018.

[76] Mojtaba Kordestani, M. Foad Samadi, and Mehrdad Saif. A New Hybrid Fault Prognosis Method for MFS Systems Based on Distributed Neural Networks and Recursive Bayesian Algorithm. *IEEE Systems Journal*, 14(4):5407–5416, December 2020.

[77] Zhanbo Cui, Bo Jing, Xiaoxuan Jiao, Yifeng Huang, and Shenglong Wang. The Integrated-Servo-Actuator Degradation Prognosis Based on the Physical Model Combined With Data-Driven Approach. *IEEE Sensors Journal*, 23(9):9370–9381, May 2023.

[78] M.S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp. A tutorial on particle filters for online nonlinear/non-gaussian bayesian tracking. *IEEE Transactions on Signal Processing*, 50(2):174–188, 2002.

[79] Tianzhi Li. Particle filter-based fatigue damage prognosis using prognostic-aided model updating. *Mechanical Systems and Signal Processing*, 211:111244, 2024.

[80] Marine Jouin, Rafael Gouriveau, Daniel Hissel, Marie-Cécile Péra, and Noureddine Zerhouni. Particle filter-based prognostics: Review, discussion and perspectives. *Mechanical Systems and Signal Processing*, 72-73:2–31, 2016.

[81] Andrea Mornacchi, George Vachtsevanos, and Giovanni Jacazio. Prognostics and Health Management of an Electro-Hydraulic Servo Actuator. In *Annual Conference of the PHM Society*, volume 7, 2015. Number: 1.

[82] Sylvain Autin, Jerome Socheleau, Andrea Dellacasa, Andrea De Martin, Giovanni Jacazio, and George Vachtsevanos. Feasibility Study of a PHM System for Electro-hydraulic Servo-actuators for Primary Flight Controls. In *Annual Conference of the PHM Society*, September 2018.

[83] Runxia Guo and Jianfei Sui. Remaining Useful Life Prognostics for the Electrohydraulic Servo Actuator Using Hellinger Distance-Based Particle Filter. *IEEE Transactions on Instrumentation and Measurement*, 69(4):1148–1158, April 2020. Conference Name: IEEE Transactions on Instrumentation and Measurement.

[84] Andrea Macaluso and Giovanni Jacazio. Prognostic and Health Management System for Fly-by-wire Electro-hydraulic Servo Actuators for Detection and Tracking of Actuator Faults. In *Procedia CIRP*, volume 59 of *Proceedings of the 5th International Conference in Through-life Engineering Services Cranfield University, 1st and 2nd November 2016*, pages 116–121, January 2017.

[85] H Kannemans and H W Jentink. A Method to Derive the Usage of Hydraulic Actuators From Flight Data. In *ICAS 2002 CONGRESS*, 2002.

[86] FC te Winkel and DJ Spiekhout. Rnlaf/f-16 loads and usage monitoring/management program. Technical Report NLR-TP-2002-309, National Aerospace Laboratory NLR, 2002.

[87] Jean-Charles Maré. Review and Analysis of the Reasons Delaying the Entry into Service of Power-by-Wire Actuators for High-Power Safety-Critical Applications. *Actuators*, 10(9):233, September 2021.

[88] Riccardo Achille, Andrea De Martin, Antonio Carlo Bertolino, Giovanni Jacazio, and Massimo Sorli. A simulation study on wear monitoring and prognosis in electro-mechanical brakes for a small passenger aircraft. *Actuators*, 15(3), 2026.

[89] Jean-Charles Maré. *Aerospace Actuators 2: Signal-by-Wire and Power-by-Wire*. John Wiley & Sons, April 2017.

[90] Christian Favre. Fly-by-wire for commercial aircraft: the airbus experience. *International Journal of Control*, 59(1):139–157, 1994.

[91] RPG Collinson. Fly-by-wire flight control. *Computing & Control Engineering Journal*, 10(4):141–152, 1999.

[92] Anton Gaile and Yuan Lue. Electro hydraulic actuation (eha) systems for primary flight control, landing gear and other type of actuation. In *2016 IEEE International Conference on Aircraft Utility Systems (AUS)*, pages 723–728, 2016.

[93] Jian Kang, Zhaohui Yuan, and Muhammad Tariq Sadiq. Numerical simulation and experimental research on flow force and pressure stability in a nozzle-flapper servo valve. *Processes*, 8(11):1404, 2020.

[94] Samuel David Iyaghigba, Fakhre Ali, and Ian K. Jennions. A Review of Diagnostic Methods for Hydraulically Powered Flight Control Actuation Systems. *Machines*, 11(2):165, January 2023.

[95] Michael Scott. *A Data-Driven approach to Prognostics and Health Management of Aircraft Structures for Predictive Maintenance in Operations and Sustainment*. PhD thesis, RMIT University, 2025.

[96] Michael J. Candon, Michael Scott, Stephan Koschel, Oleg Levinski, and Pier Marzocca. A data-driven signal processing framework for enhanced freeplay diagnostics in nextgen structural health monitoring systems. In *AIAA SCITECH 2022 Forum*, 2022.

[97] Alessandro Aimasso, Matteo D. L. Dalla Vedova, and Paolo Maggiore. Analysis of fbg sensors performances when integrated using different methods for health and structural monitoring in aerospace applications. In *2022 6th International Conference on System Reliability and Safety, ICSRS 2022*, page 138 – 144, 2022.

[98] Manuel Esperon-Miguez, Ian K. Jennions, and Philip John. Implementing ivhm on legacy aircraft: Progress towards identifying an optimal combination of technologies. In Peter W. Tse, Joseph Mathew, King Wong, Rocky Lam, and C.N. Ko, editors, *Engineering Asset Management - Systems, Professional Practices and Certification*, pages 799–812, Cham, 2015. Springer International Publishing.

[99] David E Acuna-Ureta and Marcos E Orchard. A gentle correctness verification of the theory of uncertain event prognosis to compute failure time probability. In *Annual Conference of the PHM Society*, volume 14, 2022.

[100] Francisco Roberto Jaramillo Montoya, Martín Valderrama, Vanessa Quintero, Aramis Pérez, and Marcos Orchard. Time-of-failure probability mass function computation using the first-passage-time method applied to particle filter-based prognostics. In *Annual Conference of the PHM Society*, number 1, pages 11–11, 2020.

[101] Alessandro Aimasso, Matteo D.L. Dalla Vedova, Alfredo Esposito, and Paolo Maggiore. Analysis of the fixing process of fbg optical sensors for thermomechanical monitoring of aerospace applications. In *Journal of Physics: Conference Series*, volume 2526, 2023.

[102] Alessandro Aimasso, Matteo D. L. Dalla Vedova, and Paolo Maggiore. Innovative sensor networks for massive distributed thermal measurements in space applications under different environmental testing conditions. In *2022 IEEE 9th International Workshop on Metrology for AeroSpace, MetroAeroSpace 2022 - Proceedings*, page 503 – 508, 2022.

[103] Alessandro Aimasso, Matteo Davide Lorenzo Dalla Vedova, and Paolo Maggiore. Sensitivity analysis of fbg sensors for detection of fast temperature changes. In *Journal of Physics: Conference Series*, volume 2590, 2023.

[104] Gregory W Vogl, Brian A Weiss, and Moneer Helu. A review of diagnostic and prognostic capabilities and best practices for manufacturing. *Journal of Intelligent Manufacturing*, 30:79–95, 2019.

[105] Lorenzo Ciani, Giulia Guidi, and Gabriele Patrizi. A critical comparison of alternative risk priority numbers in failure modes, effects, and criticality analysis. *IEEE Access*, 7:92398–92409, 2019.

[106] Wim JC Verhagen and Lennaert WM De Boer. Predictive maintenance for aircraft components using proportional hazard models. *Journal of Industrial Information Integration*, 12:23–30, 2018.

[107] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020.

[108] Grace Park, Vijayan Sugumaran, and Sooyong Park. A reference model for big data analytics. In *2018 9th IEEE Annual Ubiquitous Computing, Electronics & Mobile Communication Conference (UEMCON)*, pages 382–391, 2018.

[109] Andrea De Mauro, Marco Greco, and Michele Grimaldi. A formal definition of big data based on its essential features. *Library review*, 65(3):122–135, 2016.

[110] Carlos Cernuda. *On the Relevance of Preprocessing in Predictive Maintenance for Dynamic Systems*, pages 53–93. Springer International Publishing, Cham, 2019.

[111] Jing Li, Jie Wang, Guowei Zhu, Xiao Yu, Shiyu Du, and Sally Abdel. Investigation on data cleaning and analysis technology for power industry big data. In *International Conference on Big Data Analytics for Cyber-Physical System in Smart City*, pages 655–663. Springer, 2022.

[112] Fakhitah Ridzuan and Wan Mohd Nazmee Wan Zainon. A review on data cleansing methods for big data. *Procedia Computer Science*, 161:731–738, 2019.

[113] B Venkatesh and J Anuradha. A review of feature selection and its methods. *Cybern. Inf. Technol*, 19(1):3–26, 2019.

[114] Peshawa Jamal Muhammad Ali, Rezhna Hassan Faraj, Erbil Koya, Peshawa J Muhammad Ali, and Rezhna H Faraj. Data normalization and standardization: a technical report. *Mach Learn Tech Rep*, 1(1):1–6, 2014.

[115] SG Zeldam. Automated failure diagnosis in aviation maintenance using explainable artificial intelligence (xai). Master's thesis, University of Twente, 2018.

[116] Bu Yung Kosasih, Wahyu Caesarendra, Kiet Tieu, Achmad Widodo, Craig A.S. Moodie, and A. Kiet Tieu. Degradation trend estimation and prognosis of large low speed slewing bearing lifetime. In *Advances in Applied Mechanics and Materials*, volume 493 of *Applied Mechanics and Materials*, pages 343–348. Trans Tech Publications Ltd, 4 2014.

[117] M. Baghli, C. Delpha, D. Diallo, and A. Hallouche. Three-Level Inverter Fault Detection and Diagnosis Using Current-Based Statistical Analysis. In *2018 Prognostics and System Health Management Conference (PHM-Chongqing)*, pages 686–691, 2018.

[118] Sarvesh Sundaram and Abe Zeid. Smart prognostics and health management (sphm) in smart manufacturing: An interoperable framework. *Sensors*, 21(18), 2021.

[119] Udeme Ibanga Inyang, Ivan Petrunin, and Ian Jennions. Diagnosis of Multiple Faults in Rotating Machinery Using Ensemble Learning. *Sensors*, 23(2):1005, 2023.

[120] Wang Zhenya, Chen Lu, Jian Ma, Hang Yuan, and Zihan Chen. Novel method for performance degradation assessment and prediction of hydraulic servo system. *Scientia Iranica*, 22(4):1604–1615, 2015.

[121] Agnieszka Wyłomańska and Radosław Zimroz. Signal segmentation for operational regimes detection of heavy duty mining mobile machines-a statistical approach. *Diagnostyka*, 15(2):33–42, 2014.

[122] David E. Acuña-Ureta, Marcos E. Orchard, and Patrick Wheeler. Computation of time probability distributions for the occurrence of uncertain future events. *Mechanical Systems and Signal Processing*, 150:107332, 2021.

[123] Abdurrahim Bilal Özcan and Elbrus Caferov. Frequency domain analysis of f-16 aircraft in a variety of flight conditions. *International Journal of Aviation Science and Technology*, 3(01):21–34, 2022.

[124] Xuan Zhou, Michal Dziendzikowski, Krzysztof Dragan, Leiting Dong, Marco Giglio, and Claudio Sbarufatti. Generating High-Resolution Flight Parameters in Structural Digital Twins Using Deep Learning-based Upsampling. In *2023 Prognostics and Health Management Conference (PHM)*, pages 318–323, May 2023.

[125] Benedict Jux, Simon Foitzik, and Martin Doppelbauer. A standard mission profile for hybrid-electric regional aircraft based on web flight data. In *2018 IEEE International Conference on Power Electronics, Drives and Energy Systems (PEDES)*, pages 1–6, 2018.

[126] Giuseppe Vettigli. Minisom: minimalistic and numpy-based implementation of the self organizing map, 2018.

[127] Saeed Aghabozorgi, Ali Seyed Shirkhorshidi, and Teh Ying Wah. Time-series clustering – a decade review. *Information Systems*, 53:16–38, 2015.

[128] John Paparrizos, Chunwei Yang, and Guoliang Li. Bridging the gap: A decade review of time-series clustering methods. *arXiv preprint arXiv:2412.20582*, 2024.

[129] Xiaohui Huang, Yunming Ye, Liyan Xiong, Raymond YK Lau, Nan Jiang, and Shaokai Wang. Time series k-means: A new k-means type smooth subspace clustering for time series data. *Information Sciences*, 367:1–13, 2016.

[130] Xinmin Tang, Junwei Gu, Zhiyuan Shen, and Ping Chen. A flight profile clustering method combining twed with k-means algorithm for 4d trajectory prediction. In *2015 Integrated Communication, Navigation and Surveillance Conference (ICNS)*, pages S3–1–S3–9, 2015.

[131] John Paparrizos and Luis Gravano. k-shape: Efficient and accurate clustering of time series. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data*, SIGMOD '15, page 1855–1870, New York, NY, USA, 2015. Association for Computing Machinery.

[132] Ali Javed, Donna M Rizzo, Byung Suk Lee, and Robert Gramling. Somtimes: self organizing maps for time series clustering and its application to serious illness conversations. *Data Mining and Knowledge Discovery*, 38(3):813–839, 2024.

[133] John Paparrizos and Luis Gravano. k-shape: Efficient and accurate clustering of time series. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data*, pages 1855–1870, 2015.

[134] Marco Cuturi and Mathieu Blondel. Soft-DTW: a differentiable loss function for time-series. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70, pages 894–903. PMLR, 2017.

[135] Florent Forest, Mustapha Lebbah, Hanane Azzag, and Jérôme Lacaille. A survey and implementation of performance metrics for self-organized maps. *arXiv preprint arXiv:2011.05847*, 2020.

[136] Kristian Kiviluoto. Topology preservation in self-organizing maps. In *Proceedings of International Conference on Neural Networks (ICNN)*, pages 294–299. IEEE, 1996.

[137] Samuel Kaski and Krista Lagus. Comparing self-organizing maps. In Christoph von der Malsburg, Werner von Seelen, Jan C. Vorbrüggen, and Bernhard Sendhoff, editors, *International Conference on Artificial Neural Networks*, volume 1112 of *Lecture Notes in Computer Science*, pages 809–814, Berlin, Heidelberg, 1996. Springer, Springer.

[138] Hans-Ulrich Bauer and Klaus Pawelzik. Quantifying the neighborhood preservation of self-organizing feature maps. *IEEE Transactions on Neural Networks*, 3(4):570–579, 1992.

[139] Jorge E. García Bustos, Benjamin Brito Schiele, Bruno Masserano, Ricardo Salas-Espiñeira, Diego Troncoso-Kurtovic, David E. Acuña-Ureta, Francisco Jaramillo-Montoya, Marcos E. Orchard, Jorge F. Silva, and Aramis Perez. Enabling online maximum driving range prognostics in electric vehicles via uncertain event likelihood functions. *Engineering Applications of Artificial Intelligence*, 173:114449, 2026.

[140] Sidney Redner. *A Guide to First-Passage Processes*. Cambridge University Press, Cambridge, UK, 2001.

[141] Kjell A. Doksum and Arnljot Høyland. Models for variable-stress accelerated life testing experiments based on wiener processes and the inverse gaussian distribution. *Technometrics*, 34(1):74–82, February 1992.

[142] G. A. Whitmore and Fred Schenkelberg. Modelling accelerated degradation data using wiener diffusion with a time scale transformation. *Lifetime Data Analysis*, 3:27–45, 1997.

[143] Zhi-Sheng Ye, Nan Chen, and Yan Shen. A new class of wiener process models for degradation analysis. *Reliability Engineering & System Safety*, 139:58–67, July 2015.

[144] Mei-Ling Ting Lee and G. A. Whitmore. Threshold regression for survival analysis: Modeling event times by a stochastic process reaching a boundary. *Statistical Science*, 21(4):501–513, 2006.

[145] Mei-Ling Ting Lee and George A. Whitmore. First hitting time models for lifetime data. In Narayanaswamy Balakrishnan and C. R. Rao, editors, *Handbook of Statistics*, volume 23, pages 537–543. Elsevier, Amsterdam, 2003.

[146] Abhinav Saxena, Jose Celaya, Edward Balaban, Kai Goebel, Bhaskar Saha, Sankalita Saha, and Mark Schwabacher. Metrics for evaluating performance of prognostic techniques. In *2008 International Conference on Prognostics and Health Management*, pages 1–17, October 2008.

[147] Abhinav Saxena, Jose Celaya, Bhaskar Saha, Sankalita Saha, and Kai Goebel. Metrics for offline evaluation of prognostic performance. *International Journal of Prognostics and health management*, 1(1):4–23, 2010.

[148] Chenyang Lai, Piero Baraldi, and Enrico Zio. Physics-informed deep autoencoder for fault detection in new-design systems. *Mechanical Systems and Signal Processing*, 215:111420, 2024.

[149] Chenyang Lai, Piero Baraldi, and Enrico Zio. Gradient-enhanced physics-informed long short-term memory networks for stable and accurate prediction of the rul of electronic components. *Reliability Engineering System Safety*, 265:111485, 2026.

Appendix A

Appendix: Codes and Algorithms

Algorithm 1 Enhanced SOM Hyperparameter Optimization

Require: Dataset $\mathbf{X}$, patience p, threshold τ
Ensure: Optimal SOM parameters θ^*, trained SOM $\mathscr{M}^*$

Split $\mathbf{X} \to (\mathbf{X}_{\text{train}}, \mathbf{X}_{\text{val}})$
Initialize parameter grid $\mathscr{G}$
$\mathscr{S}^* \leftarrow \infty$
for $(\sigma, \alpha, T) \in \mathscr{G}$ **do**
    Create SOM $\mathscr{M}$ with parameters (σ, α)
    $(T_{\text{actual}}, \text{conv}) \leftarrow \textsc{TrainSOM}(\mathscr{M}, \mathbf{X}_{\text{train}}, T, p, \tau)$
    Compute errors: $\text{QE}_{\text{val}}, \text{TE}_{\text{val}}$
    Compute utilization rate: $\mathscr{U}_{\text{rate}}$
    $\mathscr{S} \leftarrow \text{QE}_{\text{val}} + \text{TE}_{\text{val}} + 2(1 - \mathscr{U}_{\text{rate}})$
    if $\mathscr{S} < \mathscr{S}^*$ **then**
        $\theta^* \leftarrow (\sigma, \alpha, T_{\text{actual}}, \text{conv})$
        $\mathscr{S}^* \leftarrow \mathscr{S}$, $\mathscr{M}^* \leftarrow \mathscr{M}$
    end if
end for
return θ^*, $\mathscr{M}^*$

function $\textsc{TrainSOM}(\mathscr{M}, \mathbf{X}, T_{\max}, p, \tau)$
    $\text{QE}_{\text{prev}} \leftarrow \infty$, $c \leftarrow 0$
    for $t = 1$ **to** $T_{\max}$ **do**
        Train $\mathscr{M}$ for 1 iteration
        if $t \bmod 50 = 0$ **then**
            $\text{QE}_{\text{curr}} \leftarrow$ quantization error of $\mathscr{M}$
            if $|\text{QE}_{\text{prev}} - \text{QE}_{\text{curr}}| < \tau$ **then**
                $c \leftarrow c + 1$
                if $c \geq p$ **then return** (t, True)
                end if
            else
                $c \leftarrow 0$
            end if
            $\text{QE}_{\text{prev}} \leftarrow \text{QE}_{\text{curr}}$
        end if
    end for
    return $(T_{\max}, \text{False})$
end function