Sensor Optimisation for Aircraft Health Management Systems

Burak Suslu

SCHOOL OF AEROSPACE TRANSPORT AND MANUFACTURING
Integrated Vehicle Health Management

PhD
Academic Year: 2022 - 2025

Supervisor: Prof. Ian K. Jennions
Associate Supervisor: Dr Fakhre Ali

October 2025

SCHOOL OF AEROSPACE TRANSPORT AND MANUFACTURING
Integrated Vehicle Health Management

PhD

Academic Year 2022 - 2025

Burak Suslu

Sensor Optimisation for Aircraft Health Management Systems

Supervisor: Prof. Ian K. Jennions
Associate Supervisor: Dr Fakhre Ali
October 2025

Academic integrity declaration

I declare that:

- the thesis submitted has been written by me alone.
- the thesis submitted has not been previously submitted to this university or any other.
- that all content, including primary and/or secondary data, is true to the best of my knowledge.
- that all quotations and references have been duly acknowledged according to the requirements of academic research.

understand that to knowingly submit work in violation of the above statement will be considered by examiners as academic misconduct.

Abstract

The Multi-Objective Sensor Optimisation Framework (MOSOF) views the network as an integrated system, connecting sensor selection, placement, data processing, and operational management to stakeholder-specific objectives. At the core of MOSOF is the Normalised Diagnostic Contribution Index (NDCI), which measures each sensor's worth for fault detection and isolation by merging separation ability, severity sensitivity, and informational uniqueness. Multi-objective cost functions evaluate performance, coverage, cost, and reliability, and a genetic algorithm identifies Pareto-efficient solutions.

Rigorous evaluation employs repeated nested cross-validation to prevent optimistic bias. Subsystem studies, covering the Engine, Fuel, Electrical Power System (EPS), and Environmental Control System (ECS), utilise platform "symptom vectors" to model fault propagation across subsystems. Under consistent learners, NDCI-guided suites are reliably compact and practical: NDCI provided overall better-balanced accuracy with smaller sensor sets over mRMR across three subsystems. Fuel system slightly favours mRMR on only balanced accuracy, 0.48 versus 0.53, reflecting limited sensor diversity. These results indicate that prioritising diagnostic contribution at the network level enhances detection and isolation without increasing suite size.

For platform deployment, MOSOF produces a defendable Pareto front with a knee solution consisting of 12 sensors (Engine 5, Fuel 2, EPS 2, ECS 3), achieving approximately 0.69 normalised diagnostic performance. The framework supports heterogeneous sensor types, enables multi-level data fusion, and allows certification-driven coverage or airline cost limits to be directly incorporated into the objectives.

Overall, the thesis establishes a reproducible pathway from domain-informed ranking to decision-grade trade studies. By shifting evaluation from isolated device metrics to network-level effects, MOSOF delivers compact, reliable, and economically justifiable sensor suites that improve fault coverage, isolation efficiency, and predictive maintenance across safety-critical platforms.

Keywords: Integrated Vehicle Health Management (IVHM); NDCI (Normalised Diagnostic Contribution Index); MOSOF (multi-objective sensor optimisation framework); Fault Detection and Isolation (FDI); Multi-criteria Decision-Making; Pareto Front; Multi Objective Genetic Algorithm (MOGA); Cost–Benefit Analysis; Diagnostics.

Dedicating this work to my mum, Fatma, without her last will, this PhD journey would not have been possible.
And of course, to my father, Ismail and his friend Mehmet, who became the guarantors of the funding that I received from the Ministry of Education, Türkiye.

Acknowledgements

This doctoral study was made possible through the guidance, generosity, and patience of many people and institutions.

Foremost, gratitude is recorded to Prof. Ian K. Jennions and Dr Fakhre Ali for exemplary teaching, sustained guidance, and rigorous mentorship within the School of Aerospace, Transport and Manufacturing.

Warm appreciation is extended to the IVHM Centre, Cranfield University's academic staff, research fellows, doctoral peers, and professional services for an intellectually generous environment and steadfast practical support, from technical advice and laboratory access to administrative assistance that kept the research moving.

Appreciation is also due to collaborators across the IVHM community. Valuable discussions on diagnostic reasoning and modelling with A. K. Kwakye, Dr C. M. Ezhilarasu, and L. Jia, among others, shaped the development of the Normalised Diagnostic Contribution Index and the MOSOF methodology; work published by these colleagues set an exacting benchmark that informed and motivated the present study.

Affectionate remembrance is offered for Angelia, whose passing one month prior to submission is deeply mourned. Steadfast companionship, thoughtful conversations, and gentle reminders to focus on what truly matters provided clarity and resolve during demanding periods. This thesis indeed rests, in part, on that kindness and wisdom.

Enduring encouragement from family and close friends made a sustained effort possible. Patience during extended absences, practical help at critical deadlines, and a steady belief in the value of the work provided the resilience required to complete a challenging journey.

Finally, appreciation is expressed to everyone who contributed in ways too numerous to list, by sharing code snippets, correcting figures, proofreading chapters, or posing incisive questions at the right moment. Any remaining errors and omissions remain the sole responsibility of the researcher.

Table of Contents

List of Figures

List of Tables

List of Equations

List of Abbreviations

AC	Alternating Current
ACM	Air Cycle Machine
AHMS	Aircraft Health Management System
AHP	Analytical Hierarchy Process
ANFIS	Adaptive Neuro-Fuzzy Inference System
ANOVA	Analysis of Variance
AOG	Aircraft on Ground
BBN	Bayesian Belief Network
BDE	Binary Differential Evolution
BITE	Built-In Test Equipment
BN	Bayesian Network
C-MAPSS	Commercial Modular Aero-Propulsion System Simulation
CBM	Condition-Based Maintenance
CCR	Correct Classification Rate
CDP	Compressor Discharge Pressure
CNN	Convolutional Neural Network
DBN	Dynamic Bayesian Network
DFT	Design for Testability
DL	Deep Learning
DoF	Degrees of Freedom
ECS	Environmental Control System
EPS	Electrical Power System
FA	Firefly Algorithm
FAA	Federal Aviation Administration
FAP	False Alarm Probability
FDR	Fault Detection Rate
FIM	Fisher Information Matrix
FIR	Fault Isolation Rate
FOD	Foreign Object Damage
FOHE	Fuel-Oil Heat Exchanger
FS	Fuel System
GA	Genetic Algorithm

GAN	Generative Adversarial Network
GP	Gaussian Process
HI	Health Indicator
HIL	Hardware-in-the-Loop
HPC	High-Pressure Compressor
HPT	High-Pressure Turbine
ICA	Independent Component Analysis
IG	Information Gain
IVHM	Integrated Vehicle Health Management
k-NN	k-Nearest Neighbours
KPCA	Kernel Principal Component Analysis
LPC	Low-Pressure Compressor
LPT	Low-Pressure Turbine
LSTM	Long Short-Term Memory
MCDM	Multi-Criteria Decision-Making
ML	Machine Learning
MOGA	Multi-Objective Genetic Algorithm
MOO	Multi-Objective Optimisation
MOSOF	Multi-Objective Sensor Optimisation Framework
mRMR	minimum Redundancy–Maximum Relevance
MRO	Maintenance, Repair, and Overhaul
MTBF	Mean Time Between Failures
MTTR	Mean Time to Repair
NB	Naïve Bayes
NDCI	Normalised Diagnostic Contribution Index
NDE	Non-Destructive Evaluation
OEM	Original Equipment Manufacturer
OSP	Optimal Sensor Placement
Pt_AfCDP	Pressure at the exit of bleed air duct
Pt_S24	Pressure at LPC exit (station 24)
PCA	Principal Component Analysis
PHM	Prognostics and Health Management
PHX	Primary Heat Exchanger

POD	Proper Orthogonal Decomposition
PSO	Particle Swarm Optimisation
PSV	Platform Symptom Vector
P1, .. P6	Pressure sensors 1,2,..6
RAM	Ram Air
RBF	Radial Basis Function
RMSE	Root Mean Square Error
RNN	Recurrent Neural Network
RTF	Run-to-Fail
RUL	Remaining Useful Life
S4	Systematic Sensor Selection Strategy
SAE	SAE International
SHM	Structural Health Monitoring
SHX	Secondary Heat Exchanger
SIR	Signal-to-Interference Ratio
SNR	Signal-to-Noise Ratio
STFT	Short-Time Fourier Transform
SVD	Singular Value Decomposition
T-MATS	Toolbox for Modeling and Analysis of Thermodynamic Systems
t-SNE	t-Distributed Stochastic Neighbour Embedding
Tci	Cold inlet temperature
ToC	Outlet temperature of the compressor
Tt_S5	Temperature at LPT exit (station 5)
TcoRHX	Reheater cold side outlet temperature
TCV	Temperature Control Valve
TOPSIS	Technique for Order of Preference by Similarity to Ideal Solution
TSFC	Thrust Specific Fuel Consumption
VAM	Virtual Aircraft Model
WVD	Wigner-Ville Distribution

1 Introduction

Commercial aviation operates at the intersection of high-stakes economics and safety-critical engineering. As a vital driver of the global economy, the air transport sector supports over 86 million jobs worldwide and contributes about $4.1 trillion in economic activity (approximately 3.9% of global GDP) [1, 2]. The scale of operations is immense: in 2024, airlines are expected to carry nearly 5 billion passengers on close to 39 million flights [3, 4]. Within this context, Maintenance, Repair, and Overhaul (MRO) emerges as a strategic focal point directly influencing both safety margins and financial performance. Maintenance typically accounts for 10–15% of an airline's operating costs [5], making it one of the most significant controllable expenditures after fuel and crew costs. Recent industry forecasts indicate that the MRO market, which rebounded to approximately $104 billion in 2024, will expand steadily to reach $124 billion by 2034 [6]. Even modest gains in maintenance efficiency can thus yield substantial savings for operators. Equally important, unscheduled maintenance events impose outsized costs: every minute an aircraft is grounded (AOG) incurs significant direct and indirect losses (on the order of $150,000 per day for a single plane) [7], disrupting flight schedules and eroding passenger trust. These pressures create a compelling business case for transitioning from traditional reactive maintenance to more predictive and condition-based strategies.

Figure 1-1 shows the percentage of MRO shares in total revenue and operating costs.

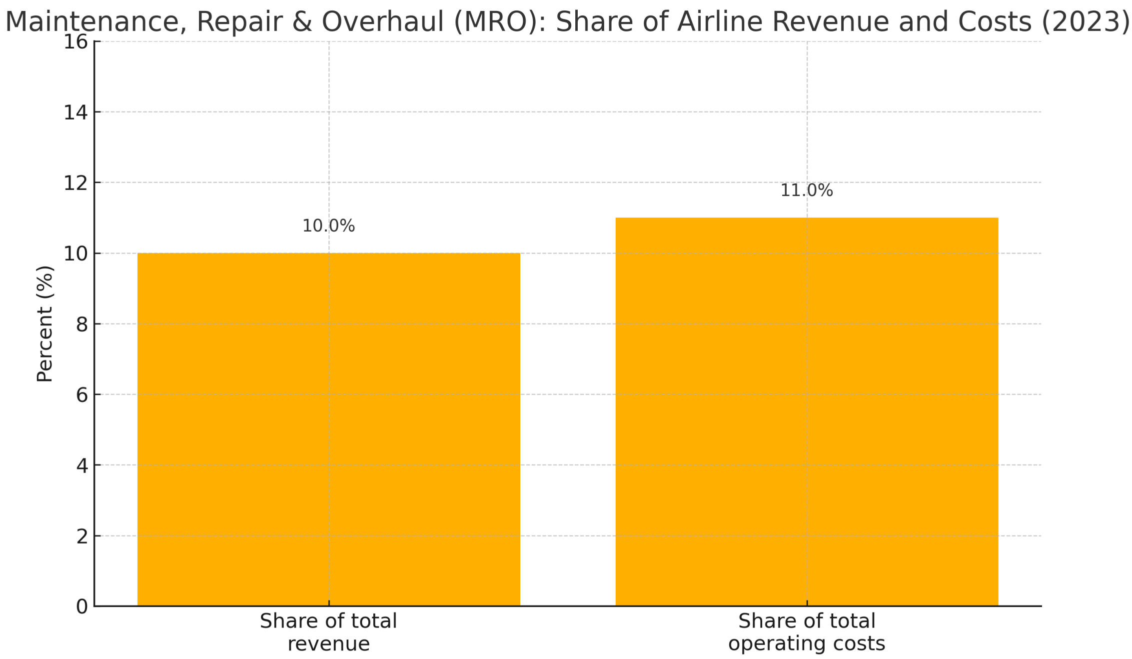

Figure 1-1 Typical share of Maintenance, Repair, and Overhaul (MRO) in airline operating expenses

At the same time, modern aircraft are being designed with increasingly integrated systems and generating unprecedented amounts of data. For example, an Airbus A350 is equipped with roughly 50,000 onboard sensors, collectively producing about 2.5 terabytes of data per day of operation [8]. This shift toward digitalisation and data-driven operations underpins advanced maintenance paradigms such as Condition-Based Maintenance (CBM) and Integrated Vehicle Health Management (IVHM). CBM/IVHM systems aim to continuously monitor aircraft health, detect incipient faults, and predict the remaining useful life (RUL) of components, enabling maintenance to be performed only when necessary, rather than at fixed intervals. In principle, such approaches promise improved aircraft availability, enhanced safety, and lower life-cycle costs. In practice, however, the real-world effectiveness of CBM/IVHM hinges critically on the quality, informativeness, and reliability of the onboard sensing infrastructure [9]. Simply put, better data from an optimally designed sensor network translates into more accurate diagnostics and prognostics. This reality motivates the central focus of this thesis: a rigorous, systematic approach to optimising the sensor networks that support modern Aircraft Health Management Systems (AHMS).

This introduction establishes the motivation and context for the research, defining its main challenges and outlining the thesis aim, objectives, and structure. It begins by examining the imperative of sensor optimisation in aviation maintenance (Section 1.1), then discusses the inherently multi-objective nature of the problem (1.1.1) and the challenges posed by complex, interconnected systems (1.1.2). The second subsection identifies a methodological gap in current sensor network design practices (Section 1.2) that this research aims to address and system health indicator(1.2.1). The subsequent sections of this chapter will articulate the research aim and objectives (Section 1.3), provide an overview of the methodological approach (1.4), outline the thesis structure (1.5), and the list of publications (1.6) that have emanated from this research.

1.1 The Imperative of Sensor Optimisation

The cornerstone of any modern AHMS is the network of sensors distributed throughout the aircraft. These sensors provide the raw data needed to monitor the health and performance of countless critical components. However, the effectiveness of an AHMS is not merely a function of the quantity of sensors deployed; it fundamentally depends on their optimal configuration and deployment. Simply adding more sensors is rarely a cure – redundant or poorly placed sensors can inflate system cost, weight, and power consumption while failing to capture the most salient fault indicators. In fact, decades of research in fault diagnosis and system monitoring have shown that sensor selection and placement decisions significantly influence the ability to detect, isolate, and identify failures [9, 10]. *Sensor optimisation*, in the context of this thesis, is defined as the strategic process of:

- Selecting the appropriate sensing modalities and sensor quality grades for each measurement point (ensuring that each chosen sensor type is fit for purpose).
- Determining optimal sensor placement and network topology, so that sensors are located where they capture the most informative signals for fault detection while minimising redundancy.
- Defining operational policies (such as sampling rates, thresholds, and adaptive sensing modes) to balance information quality against resource consumption (power, bandwidth, data storage).
- Improving data quality and management, including signal processing, data fusion, and health monitoring algorithms that maximise the utility of the sensor data collected.

An optimised sensor network ensures that the data fed into the health management system is of the highest fidelity and relevance, enabling accurate assessment of system state, timely fault detection, and reliable prognostics. In practical terms, optimisation means getting the “right” data from the “right” locations at the “right” times. A well-chosen suite of sensors will capture incipient fault signatures that would be missed by a naively designed or legacy sensor configuration. Conversely, a suboptimal sensor setup – for example, using too few sensors or monitoring the wrong parameters – can leave critical failure modes unmonitored. In contrast, an overly dense sensor network can overwhelm data handling systems and maintenance crews with irrelevant information. Modern aircraft already produce enormous datasets (recall the A350 example), so there is a premium on selecting smart sensors rather than relying on brute-force instrumentation [8]. In summary, sensor optimisation is imperative because it directly governs the quality of information available for decision-making in maintenance and safety: an aircraft can only be as “healthy” and failure aware as the sensors that are watching over it.

1.1.1 Multi-Objective Nature of the Problem

Designing an optimal sensor network for a complex aerospace system is inherently a multi-objective optimisation problem. Engineers must navigate a landscape of competing (often conflicting) design goals and constraints [9]. For instance, maximising diagnostic coverage and accuracy may call for a dense network of high-precision sensors to capture every possible fault signature. However, such a maximalist approach would inevitably increase the aircraft's cost, add weight (which in turn impacts fuel burn and performance), and consume more power – all of which are tightly constrained parameters in aerospace design [9]. There is thus a fundamental tension between performance objectives (e.g., fault detectability, isolation resolution, prognostic horizon) and cost/constraint objectives (e.g., acquisition and maintenance costs, weight, volume, power usage).

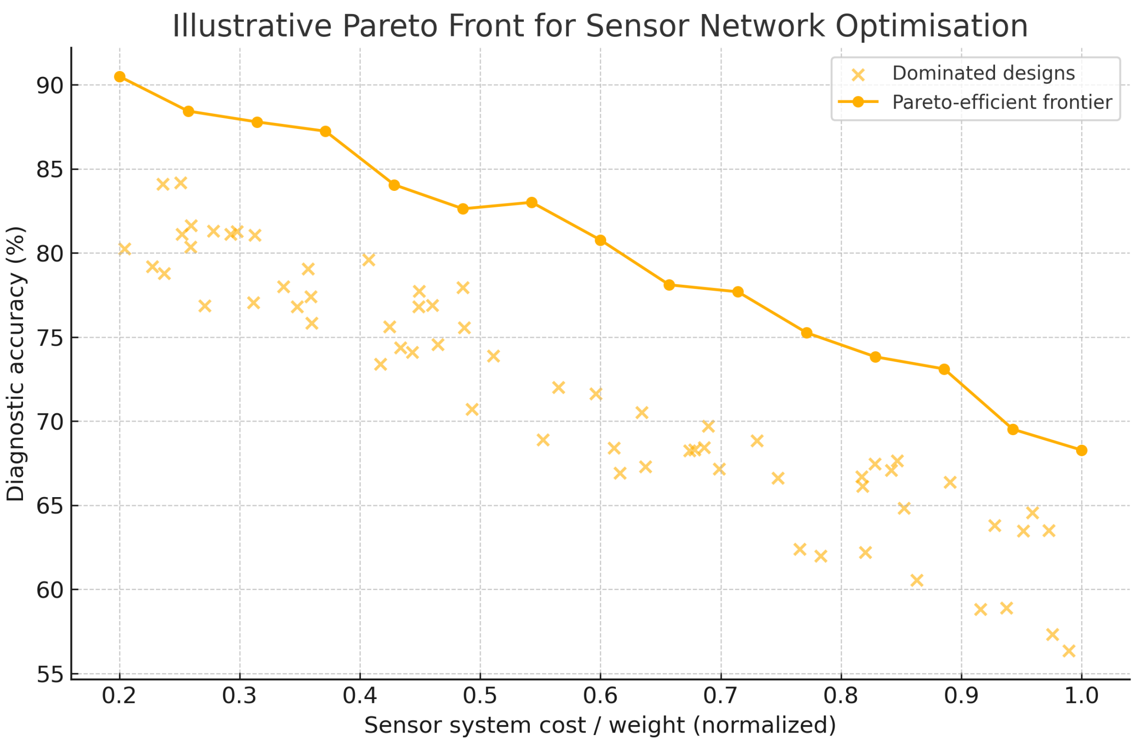

Figure 1-2 Conceptual illustration of the Pareto frontier in multi-objective sensor optimisation, depicting the fundamental trade-off between diagnostic performance and system-level costs

This conflict creates a classic engineering trade-off that can be visualised as a Pareto frontier (Figure 1-2 conceptually illustrates this idea). For example, minimising system-level costs by reducing the number of sensors or using only low-cost devices could jeopardise the reliability and effectiveness of the AHMS, potentially allowing critical failures to go undetected or unanalysed. On the other hand, pushing for the highest possible fault coverage with numerous high-grade sensors would yield diminishing returns and unsustainable costs. The multi-objective nature of sensor network design generates a discrete Pareto frontier, illustrating the non-convex trade-off between diagnostic utility and system constraints. As depicted in Figure 1-2, this frontier exhibits three distinct topological features. Initially, there is an ascent phase, where core thermodynamic sensors provide maximum mutual information, yielding exponential diagnostic gains for minimal initial investment. This is followed by a knee, representing the point of diminishing marginal returns due to signal correlation. Finally, the curve reaches a saturation plateau, where diagnostic gain is limited by the fundamental observability of the physical system and additional sensors merely introduce redundant noise. Furthermore, any viable engineering solution must reside above a strict certification floor, which represents the minimum fault detection coverage mandated by airworthiness safety standards.

The challenge for the designer is to find balanced solutions that offer the best compromise – in other words, to identify sensor configurations that are Pareto-optimal (non-dominated by any other solution in all objectives). Modern optimisation techniques provide the tools to explore this trade space systematically. In particular, methods from multi-objective optimisation (MOO), such as evolutionary algorithms (e.g. NSGA-II) and Pareto ranking strategies, are well-suited to this task. These algorithms enable the simultaneous consideration of multiple criteria and can efficiently approximate the set of optimal trade-off solutions [11,12]. By applying such methods to sensor network design, a Pareto frontier of solutions that span, for example, the spectrum from low-cost/low-performance configurations to high-cost/high-performance ones can be obtained. Decision-makers (aircraft manufacturers, airlines, MRO providers) can then select an acceptable point on this frontier based on their priorities or constraints.

It is worth noting that the multi-objective nature extends across different stakeholders. An aircraft manufacturer (OEM) may prioritise design criteria differently from an airline operator or a maintenance service provider. For instance, an OEM might focus on certifiability and warranty costs, an airline on dispatch reliability and life-cycle cost, and an MRO on ease of maintenance and turnaround time. A comprehensive sensor optimisation framework must accommodate these perspectives, ensuring that the chosen solution is well-balanced and effective.

In the context of this sensor optimisation analysis, specific assumptions are made regarding the temporal horizons and primary economic drivers of each stakeholder, which inherently split the optimisation surface. It is assumed that OEMs are primarily driven by unit production cost to ensure competitive initial aircraft pricing, alongside strict adherence to certification requirements. Their evaluation horizon is inherently front-loaded. Conversely, Airline operators are assumed to evaluate sensor utility over a 20-year lifecycle horizon; thus, their primary drivers are direct maintenance costs, fleet availability, and long-term dispatch reliability. MRO providers are assumed to focus heavily on operational maintainability, requiring sensor suites that enable rapid fault isolation to reduce aircraft-on-ground turnaround times. Consequently, a foundational assumption of the MOSOF analysis is that there is no single, universal 'optimal' point; rather, optimality is strictly relative to the temporal horizon and cost-penalty assumptions of the evaluating stakeholder.

Modern aircraft are quintessential examples of complex systems. They comprise numerous highly interconnected subsystems, including propulsion (engines), electrical power generation and distribution, fuel management, environmental control (air conditioning and pressurisation), hydraulics, avionics, and many more, all operating in concert to ensure the aircraft’s functionality and safety [13]. These subsystems interact in intricate ways: a fault in one system can propagate or manifest as anomalies in another, seemingly unrelated system. For example, a sticking engine bleed-valve not only starves the ECS of conditioned air but simultaneously shifts combustion dynamics measurable in exhaust gas temperature and N1, and — through ACM torque changes — manifests as an abnormal electrical load in the EPS. Such failure propagation is not just a theoretical possibility; real incidents evidence it and are a key challenge identified in vehicle-level health management research [14,15].

Traditional maintenance paradigms and diagnostic approaches, which tend to treat each subsystem in isolation, struggle with this complexity. Moreover, because aircraft operate under widely varying conditions (different flight phases, power settings, altitudes, and environmental conditions), sensor readings and fault signatures are context-dependent [16, 17]. A practical sensor optimisation framework must therefore embrace a systems-of-systems perspective, and a system-aware optimisation that considers the aircraft's integrated nature [18]. By doing so, the resulting sensor network can more reliably distinguish genuine faults, from cascade effects or benign fluctuations. Recent IVHM research efforts underscore the importance of this perspective. For instance, Ezhilarasu *et al.* propose an integrated diagnostic reasoning framework at the aircraft level, where multiple subsystem health monitors are combined with a top-level reasoner to identify propagated faults [15,19]. Their work highlights that current industry practice, which focuses on automated fault isolation mainly at the subsystem level, leaves a gap in handling cascading faults, a gap that integrated, vehicle-level health management aims to fill [15].

In summary, multi-objective optimisation is not just a mathematical formalism here, but a necessity: it provides a rigorous way to balance diagnostic performance against the practical costs and penalties associated with an extensive sensor network [9]. Embracing this approach enables the framework to overcome the limitations of ad-hoc or single-objective design methods and systematically explore trade-offs, ultimately identifying solutions that deliver strong health monitoring capabilities without violating real-world constraints on weight, power, or budget.

1.1.2 Sensor Optimisation as a Foundation for Advanced Health Management

The ongoing evolution from scheduled, interval-based maintenance to condition-based and predictive maintenance represents a significant shift in aviation maintenance. IVHM systems utilise real-time sensor data to continuously assess aircraft health, detect incipient faults, and predict the remaining useful life of components [17, 20]. These capabilities enable precise maintenance actions to be performed when needed, rather than reacting to a failure or adhering to conservative fixed schedules. However, a precondition for the success of such advanced health management strategies is a well-designed sensor infrastructure. In essence, an optimised sensor network is the bedrock upon which IVHM capabilities are built. High-quality, well-placed sensors provide the observability required for sophisticated diagnostic algorithms (including model-based reasoning and machine learning) to identify subtle deviations from normal behaviour. For prognostics, capturing the right indicators (e.g. vibration signatures, oil debris counts, temperature trends) is critical for forecasting component degradation and failure timelines with confidence.

By ensuring that the “right” data is captured from the “right” locations, sensor optimisation greatly enhances diagnostic resolution and prognostic fidelity. For example, the early stages of a fault often produce only minor perturbations in sensor readings; without a sensitive sensor in the correct location, these early warnings will be missed. An optimised suite might include higher-fidelity sensors or redundant sensing for parameters that are strongly indicative of failure modes, thereby improving the sensitivity and specificity of fault detection. Research has shown that using domain-informed sensor selection can substantially improve the ability to isolate faults compared to generic feature selection. Effectively, better sensors lead to earlier and more reliable detection of problems [10]. In turn, this enables longer lead times for maintenance crews to react, more accurate prediction of RUL, and overall, more predictable and efficient maintenance operations.

Furthermore, an optimised sensor network contributes directly to safety and regulatory compliance. In aircraft design, safety assessments (such as those guided by ARP4761) identify hazardous failure modes and the mitigations required to handle them. Sensors often serve as mitigation means, providing indications or warnings that enable pilots or automated systems to take corrective action before a fault leads to an unsafe condition. A systematic sensor optimisation process can explicitly account for these safety requirements by allocating sensors to monitor critical failure paths. For instance, if loss of pressure in a hydraulic circuit is classified as a hazardous event, the optimisation framework can ensure pressure sensors (or redundant sensors) are placed appropriately to detect it. By doing so, the approach supports compliance with stringent certification guidelines and fault tolerance criteria. Regulatory bodies, such as the FAA, have traditionally relied on prescriptive maintenance and inspections to ensure safety, and have only gradually accepted condition-based approaches as evidence of continued airworthiness [21, 22]. However, there is recognition in the industry that validated health monitoring systems can earn “maintenance credits” – allowing the extension of inspection intervals or other relief, provided it can be demonstrated that an IVHM system meets the necessary reliability and safety standards [22, 23]. Notably, guidelines such as FAA AC 29-2C MG-15 (for rotorcraft) already provide a pathway for airworthiness approval of health usage monitoring systems, with the caveat that the probability of missing a hazardous condition must be extremely low [24,25]. In this light, sensor optimisation is not only an engineering endeavour but also part of the regulatory and certification strategy for new aircraft systems. An optimised sensor network can be designed to meet Design Assurance requirements (as per standards such as SAE ARP4754A/B for complex systems) and to demonstrate the required detection coverage and redundancy for safety-critical parameters, thereby upholding the aerospace industry’s uncompromising safety standards.

In summary, advanced health management capabilities (like real-time diagnostics and prognostics) are only as effective as the sensor data that feeds them. Optimising the sensor network is therefore a foundational step – it maximises the value of available information, ensures compliance with safety requirements, and amplifies the benefits of IVHM (reduced downtime, lower maintenance costs, and higher reliability). This thesis builds on that premise, aiming to show that a rigorous optimisation approach to sensor configuration can indeed unlock significant gains in an aircraft's health management performance.

1.2 The Research Problem: A Methodological Gap in Sensor Network Design

Despite the critical role of sensor networks in modern aircraft health management, the methodologies commonly used to design these networks have not kept pace with the growing complexity of the systems they monitor. This has created a significant gap between the theoretical potential of IVHM and its practical implementation. In particular, current practice is limited by three core deficiencies in sensor network design methodology. These deficiencies are interrelated and mutually reinforcing, resulting in a systematic shortcoming that this research seeks to address.

First, sensor selection and placement are inherently a multi-objective problem (as discussed in Section 1.1.2); yet, in practice, they are often approached with a single-objective or ad-hoc mindset. In industry, it is not uncommon to find that sensor decisions are driven primarily by cost or engineering convenience, with other considerations often treated as afterthoughts. Such simplistic approaches cannot systematically balance the competing goals that engineers and stakeholders face, detailed in Table 1-1. For example, maximising diagnostic performance and coverage would push toward more sensors and more sophisticated sensing, whereas minimising cost, weight, and power goes in the opposite direction. If treated in a single-objective manner (for example, minimising cost under a rudimentary performance constraint), the design process may overlook configurations that offer a significantly better trade-off. The trade space is rich and complex, involving objectives such as fault detectability, isolation accuracy, sensor reliability, weight impact, life-cycle cost, and even considerations like sensor commonality or maintenance overhead [9].

Navigating this trade-space requires multi-objective optimisation techniques, yet many traditional design processes lack such formal methods – instead relying on rules of thumb or incremental tweaks to legacy sensor layouts. The result is that Pareto-optimal solutions (which represent the best compromises among criteria) remain unexplored or unrecognised. In sum, there is a methodological gap in systematically treating sensor network design as the multi-objective optimisation problem that it truly is.

Table 1-1 Key Trade-offs in Sensor Network Design for Aerospace Stakeholders

Design Criterion	Original Equipment Manufacturer (OEM)	Airline / Operator	Maintenance, Repair & Overhaul (MRO) Provider
Diagnostic Performance	Medium: Must meet certification requirements for fault detection.	High: Directly impacts operational reliability and prevents costly unscheduled maintenance.	High: Enables rapid and accurate fault isolation, reducing aircraft-on-ground time.
Acquisition Cost	High: A primary driver of unit production cost and market competitiveness.	Medium: Viewed as part of the total cost of ownership; willing to invest for long-term ROI.	Low: Not a direct cost but influences the affordability of diagnostic tools and training.
Weight / Fuel Burn	High: A key performance parameter and selling point for the aircraft.	High: Directly translates to recurring operational costs over the aircraft's lifecycle.	Low: Indirect impact through part logistics.
System Reliability (MTBF)	High: Affects warranty costs and brand reputation for reliability.	High: Directly impacts fleet availability and schedule integrity.	Medium: Affects the reliability of the diagnostic system itself, influencing trust and efficiency.

Certifiability	High: Non-negotiable requirement for aircraft type certification (e.g., compliance with ARP4754B).	High: A prerequisite for operating the aircraft.	N/A
Maintainability	Medium: Influences the design of maintenance procedures and manuals.	High: Directly impacts the time and cost required to perform maintenance tasks.	High: The primary driver of MRO business efficiency and turnaround times.

Table 1-1 introduces the perspectives of different stakeholders on sensor selection. These categories are analytical archetypes, not prescriptive rules; they are used to parameterise the multi-objective optimisation process. For example, OEMs may emphasise certification, integration burden and warranty exposure, Airlines may prioritise operational reliability and lifecycle cost, and MROs may focus on maintainability and fault-isolation speed. Framing sensor selection within these archetypes enables MOSOF to align diagnostic performance, cost and reliability with stakeholder-specific objectives.

Second, the selection of informative sensors is often guided by generic, data-driven metrics that are "diagnostically blind" – meaning they are agnostic to the underlying physics and failure modes of the system. In this research an informative sensor is one whose measurements provide distinct, fault-relevant and non-redundant information that improves fault detection, isolation or severity tracking beyond what is already provided by the remaining sensor suite. Informative sensors maximise diagnostic metrics such as detection rate, fault isolation rate and severity-tracking accuracy while reducing uncertainty and information redundancy. A prime example is the use of statistical feature selection algorithms like minimum Redundancy–Maximum Relevance (mRMR) for choosing sensors or signals. mRMR (and similar techniques) rank candidate sensors or features based on correlation with fault events (relevance) while penalising redundancy among selected features. While powerful in many machine learning applications, such methods have fundamental limitations when applied to engineering diagnostics. They do not inherently understand the system's causal structure or the relative *criticality* of different faults. A sensor that is statistically less informative in each dataset might be essential for detecting a rare-but-catastrophic failure mode; conversely, a sensor that correlates well with a common fault might be selected multiple times in different forms (providing redundant information). In other words, generic algorithms can inadvertently discard sensors that are crucial for robust fault isolation or overweight sensors that contribute little new insight. This shortcoming has prompted calls for domain-specific metrics that quantify a sensor’s actual *diagnostic value*, rather than relying purely on statistical relevance [10,16]. For example, one could use *fault diagnosability indices* derived from system models – metrics that check how well a given sensor set distinguishes between all relevant fault scenarios. Notably, researchers at NASA and elsewhere have demonstrated model-based sensor placement techniques that explicitly aim for diagnosability. By ensuring all target faults can be uniquely isolated, and doing so with the minimum number of sensors, these methods directly tie sensor selection to diagnostic performance [10]. The absence of such diagnosability-driven design in common practice means current sensor networks may not be truly optimal for fault diagnosis – they are often designed with a “one-size-fits-all” data-driven approach that might miss subtle needs of safety-critical fault detection. This represents a second facet of the methodological gap.

Third, traditional design processes often suffer from a “system-of-systems blind spot”. The modern aircraft systems are deeply interconnected, but design responsibilities and analyses are frequently partitioned by subsystem. Sensor placement may be considered separately for the engine, fuel system, electrical system, and other components, each under a different team or using separate criteria. This siloed approach inevitably fails to address fault propagation across subsystem boundaries – the very issue that integrated health management aims to address. Optimising a sensor suite in isolation for each subsystem can lead to diagnostic blind spots at the interfaces, where a fault in one system can go undetected or be misinterpreted because its effects appear in another system’s sensors that were not designed to capture it. In practice, this could mean that specific cross-system failure modes (e.g., cascading failures involving multiple subsystems) may not be detected until a secondary effect becomes severe. There is growing recognition that a vehicle-level perspective is needed. As one study notes, current diagnostic practice focuses on subsystem-level fault isolation, and it is precisely the integrated reasoning across subsystems that is lacking, a gap that IVHM aims to fill [15]. Thus, a third deficiency is the lack of a system-of-systems approach in sensor network design. Without addressing this, even a multi-objective, diagnostically informed sensor optimisation could fall short, because it might optimise each piece but not the whole.

Significantly, these three deficiencies are deeply intertwined rather than independent. A narrow subsystem focus (the third issue) makes it difficult to correctly assess a sensor’s actual value (contributing to the second issue), because some sensors may only prove their worth when cross-system interactions are considered. Conversely, the lack of a proper value metric (second issue) encourages designers to fall back on simplistic measures like mRMR or arbitrary rules, which in turn makes the multi-objective trade-offs (first issue) harder to navigate – it is challenging to balance cost vs. benefit when “benefit” is not accurately quantified. The complexity of the trade-space might then lead practitioners to default to single-objective thinking or to copy previous designs. In effect, these deficiencies create a vicious cycle of methodological inadequacy that perpetuates suboptimal sensor network designs.

In summary, there is a clear and compelling research gap: the absence of a systematic, multi-objective, and system-aware framework for sensor optimisation, one that is guided by domain-specific diagnostic value rather than generic proxies. Filling this gap is the central aim of this thesis. The research posits that by addressing all three aspects, employing proper multi-objective optimisation, using diagnosability-based metrics, and considering the entire system of systems. A sensor network designs that markedly improve aircraft health management outcomes can be developed. The forthcoming chapters will detail the development of a framework called the Multi-Objective Sensor

Optimisation Framework (MOSOF) and demonstrate its application in aerospace case studies. By doing so, this work aims to advance the state of the art in IVHM sensor network design, bridging the methodological gap identified above and contributing to safer and more efficient aircraft operations.

1.2.1 System health indicators

IVHM provides a unified framework for assessing the condition of complex aerospace systems by integrating sensing, data processing, and decision-support functions across multiple subsystems. Rather than treating subsystems independently, IVHM adopts a layered architecture in which raw sensor data are transformed into progressively higher levels of abstraction, including condition monitoring, fault detection, health assessment, and prognostics [11], [12], [23]. Within this architecture, health indicators (HIs) serve as the critical interface between low-level measurements and higher-level diagnostic reasoning processes.

A HI can be defined as a low-dimensional representation derived from sensor data that reflects the underlying degradation state of a component or subsystem. In contrast to raw measurements or simple residuals, HIs encapsulate diagnostically meaningful information in a form suitable for fault detection, isolation, and degradation tracking. From a system-of-systems perspective, an HI is not merely a feature, but a structured representation that must remain interpretable and comparable across subsystems operating under varying conditions [11], [13].

For an HI to be effective in an IVHM context, it must satisfy several fundamental properties. First, it should exhibit strong separation capability, enabling clear discrimination between healthy and faulty states. Second, it must demonstrate severity sensitivity, providing a consistent and preferably monotonic response to increasing fault magnitude. Third, it should ensure uniqueness, contributing information that is not redundant with other indicators. These characteristics are consistent with widely recognised HI evaluation criteria such as monotonicity, trendability, and prognosability, which are extensively discussed in the prognostics and health management literature [26], [27]. From an information-theoretic perspective, these requirements align with the need to maximise relevance while minimising redundancy in feature selection [17], [18].

Historically, aircraft health monitoring relied on binary threshold-based diagnostics, where deviations from nominal behaviour triggered discrete fault flags. While effective for detecting abrupt failures, such approaches are inherently limited in their ability to capture gradual degradation processes and often fail to provide sufficient early warning for predictive maintenance. Modern IVHM systems therefore rely increasingly on continuous health indicators, which track the evolution of system condition over time and enable condition-based maintenance strategies [26], [28]. This transition is particularly important in complex, interconnected systems, where fault propagation across subsystems may only become observable at certain severity levels, making binary indicators insufficient for robust diagnostics [13].

The extraction of robust HIs from high-dimensional and noisy sensor data remains a central challenge. Three principal classes of methodologies can be identified. Physics-based approaches derive indicators from first-principles models, typically through residual generation or estimation of physical health parameters, offering strong interpretability and alignment with engineering models. Data-driven approaches, including statistical and machine learning techniques, extract latent features directly from data and are capable of capturing complex and nonlinear relationships [14], [15], [29]. However, they often depend on large datasets and may lack interpretability. As a result, hybrid approaches, which combine physical insight with data-driven methods, are increasingly adopted to balance robustness, scalability, and interpretability in aerospace applications [13], [30].

Crucially, the effectiveness of any HI is fundamentally constrained by the sensor network from which it is derived. The diagnostic value of an indicator depends not only on the extraction methodology but also on the availability, placement, and quality of the underlying measurements. This establishes a direct link between HI design and sensor optimisation. In complex system-of-systems environments, where fault propagation spans multiple subsystems, sensor selection must therefore be guided by the diagnostic contribution of each sensor rather than purely statistical relevance. Approaches based solely on generic feature selection metrics, such as mutual information or mRMR, may fail to capture the true diagnostic importance of sensors, particularly for rare or safety-critical faults [17], [18].

To address this limitation, recent research has emphasised the use of diagnosability-driven sensor selection, where sensor configurations are evaluated based on their ability to distinguish between fault scenarios and improve fault isolation ability [19], [20]. This perspective highlights the need for domain-specific metrics that explicitly quantify diagnostic usefulness in terms of separation, sensitivity, and uniqueness. Building on this foundation, this thesis adopts the NDCI as a quantitative measure of sensor informativeness and integrates it within the MOSOF [31], [32]. This approach enables the identification of sensor configurations that maximise diagnostic capability while satisfying practical engineering constraints such as cost, weight, and reliability. By explicitly linking health indicator quality to sensor selection, this work addresses a key limitation in existing IVHM methodologies, where sensor networks are often designed without direct consideration of the diagnostic content they enable.

1.3 Aim and Objectives

The overarching aim of this research is to establish and validate a novel, systematic framework for the multi-objective optimisation of sensor networks to enhance the health management of complex aerospace systems.

This framework, termed the Multi-Objective Sensor Optimisation Framework (MOSOF), is designed to provide a structured, system-aware methodology that addresses the critical trade-offs inherent in developing diagnostic and prognostic systems for modern civil aircraft. The ultimate goal is to demonstrate that a rigorous optimisation approach can yield sensor configurations that lead to quantifiable improvements in diagnostic performance, operational reliability, and maintenance efficiency.

The following research objectives delineated are the product of an iterative methodological evolution, shaped by three distinct phases of inquiry as the project progressed. Initially, the research adopted a component-centric approach, focusing on the optimisation of isolated subsystems. However, the identification of complex fault propagation pathways, such as cross-system effects from the Engine to the ECS, demonstrated that localised optimisation was insufficient, prompting a necessary expansion to a 'System-of-Systems' perspective. Subsequently, a metric-driven phase revealed the inadequacy of standard statistical feature selection algorithms for evaluating diagnostic utility, particularly their inability to capture continuous fault severity. This methodological gap necessitated the development of a novel, physics-aware diagnostic metric, the NDCI. Finally, the project culminated in a holistic integration phase. Acknowledging that technical optimisation must remain tethered to economic and operational realities, the overarching aim transitioned from singular performance maximisation to multi-objective Pareto optimisation, explicitly incorporating diverse stakeholder constraints.

To achieve this aim, the following specific objectives have been defined:

1. To establish the theoretical foundation of this research by systematically reviewing and synthesising the state-of-the-art in sensor optimisation methodologies, multi-objective optimisation algorithms, and diagnostic performance metrics. This critical analysis will identify the current research gaps and define the specific requirements for the proposed framework.
2. To design and formulate the novel Multi-Objective Sensor Optimisation Framework (MOSOF). This framework will be architected to:
    - Incorporate versatile, multi-objective cost functions that explicitly balance competing criteria such as diagnostic performance against system-level constraints like sensor cost and weight.
    - Provide a structured, repeatable methodology for the integrated tasks of sensor selection, placement, and operational configuration.
    - Embody a system-aware philosophy that accounts for the complex interdependencies and fault propagation paths within modern aircraft.
3. To develop a practical diagnostic metric, the Normalised Diagnostic Contribution Index (NDCI), and integrate it within the MOSOF architecture. This objective focuses on creating a quantifiable measure of a sensor's contribution to system-level diagnostics, with an initial application and proof of concept on a critical aircraft Environmental

Control System (ECS).

4. To demonstrate the generality and scalability of the MOSOF methodology by applying it across a diverse set of four key interlinked aircraft subsystems: the Electrical Power System (EPS), the Fuel System (FS), the Engine, and the Environmental Control System (ECS).
5. To rigorously validate the performance of the fully developed NDCI-MOSOF framework through a series of realistic use-case scenarios that simulate a range of aircraft operational conditions and fault signatures. The effectiveness of the optimised sensor networks will be quantitatively evaluated against established diagnostic performance metrics to demonstrate a tangible improvement over baseline configurations.
6. To critically assess the contributions and limitations of the research and to chart a course for future inquiry. This includes proposing specific directions for enhancing the MOSOF framework, exploring new application domains, and considering the impact of emerging trends in sensor technology and artificial intelligence on aircraft health management.

These objectives form a coherent and logical research plan. Each objective represents a distinct phase of the work, building upon the last to systematically progress from theoretical review to framework development, practical application, and rigorous validation. The successful completion of these steps will fulfil the overarching research aim, delivering a significant contribution to the field of aircraft health management.

In pursuit of this aim, the research is guided by the following key research questions (RQ):

RQ1: How can sensor networks for aircraft health management be optimised in a multi-objective manner that balances diagnostic performance with cost, weight, and other system-level constraints?

RQ2: What quantitative index can be developed to measure each sensor’s contribution to fault diagnosis, and how can this index be used to guide the sensor selection process?

RQ3: To what extent can a sensor optimisation framework, augmented by such a diagnostic metric, improve fault detection and isolation across multiple, interconnected aircraft subsystems compared to existing sensor selection approaches?

Collectively, achieving these objectives and answering these questions will accomplish the research aim and deliver a significant advancement in the field of aircraft health management. It should be noted that these objectives represent the refined culmination of an iterative process. The research began with broad multi-objective optimisation using existing feature-selection metrics; it was the identification of mRMR's insensitivity to fault severity that motivated the development of the NDCI as a dedicated diagnostic contribution metric.

1.4 Methodology

To achieve the aim and objectives outlined in Section 1.3, this research employs a structured, five-phase methodology. This phased approach ensures a systematic and rigorous investigation, progressing from the acquisition of foundational knowledge to the development, implementation, validation, and conclusive evaluation of the framework. Each phase builds upon the previous one, creating a coherent and progressive research trajectory as illustrated in Figure 1-3 and 1-4.

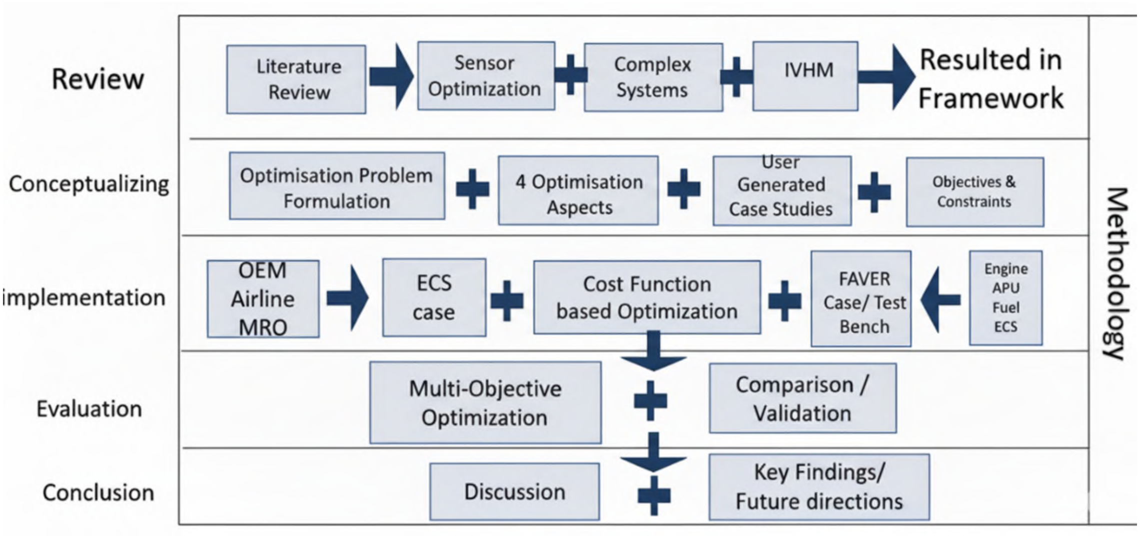

Figure 1-3 The initial research trajectory

Figure 1-4 The five-phase research methodology, showing the progression from literature review to final evaluation and the key deliverables of each phase

The five phases are detailed in the following subsections.

1.4.1 Phase 1: Literature Review

The initial phase of this research is dedicated to a comprehensive and critical review of existing literature relevant to sensor optimisation, complex systems health management, and diagnostic index development. This phase serves as the bedrock for the entire research endeavour, providing a strong theoretical foundation and identifying key research gaps. The literature review encompasses the following key areas:

- Sensor Optimisation Methodologies: This will involve a detailed examination of various sensor optimisation techniques applicable to complex systems. The review will explore different approaches to sensor selection, placement strategies, data acquisition optimisation, and sensor network management. It will analyse both traditional methods and more recent advancements in the field, including computational intelligence-based optimisation techniques.
- Multi-Objective Optimisation in Engineering Design: Given the inherent multi-objective nature of sensor optimisation, this review area will focus on multi-objective optimisation methodologies. It will investigate various techniques for handling conflicting objectives, including Pareto optimality, weighted sum methods, epsilon-constraint methods, and evolutionary algorithms. The suitability of different multi-objective optimisation approaches for sensor network design will be critically assessed.
- Diagnostic Indices for System Health Assessment: This component of the literature review will focus on the development and application of diagnostic indices in complex systems, particularly within the context of health management. It will explore different types of diagnostic indices, their mathematical formulations, their sensitivity to various fault conditions, and their effectiveness in providing reliable assessments of system health. Specific attention will be paid to indices used in aerospace and related engineering domains.
- Applications of Sensor Optimisation in Aerospace and Similar Complex Systems: To contextualise the research, this review will investigate existing applications of sensor optimisation within the aerospace industry and other domains characterised by complex systems, such as automotive, energy, and manufacturing. This will identify successful strategies, challenges encountered, and lessons learned from prior implementations, informing the design of the MOSOF framework.
- Gap Analysis: A critical outcome of the literature review will be the identification of gaps in current research and practice. This gap analysis will pinpoint areas where existing methodologies fall short in addressing the specific challenges of sensor optimisation for complex aircraft systems, thereby justifying the need for the proposed MOSOF framework and guiding its development.

The literature review is conducted using a range of academic databases, including IEEE Xplore, Scopus, Web of Science, and Google Scholar. Relevant journals and conference proceedings in the fields of aerospace engineering, sensor technology, systems engineering, and optimisation will be systematically searched. The findings of this phase are synthesised and presented in Chapter 2, "Understanding the Role of Sensor Optimisation in Complex Systems," providing a comprehensive background and context for the subsequent phases of the research.

1.4.2 Phase 2: Use Case Conceptualisation

Phase 2 focuses on conceptualising relevant use cases to guide the development and evaluation of the MOSOF framework. These use cases will be grounded in realistic scenarios within civil aircraft operations and maintenance, ensuring the practical relevance and applicability of the research outcomes. This phase involves:

- Identification of Representative Aircraft Subsystems: Based on the literature review and industry relevance, specific aircraft subsystems will be selected as representative examples for applying and testing the MOSOF framework. These subsystems will be chosen to reflect varying levels of complexity and criticality within the aircraft, ensuring a comprehensive evaluation of the framework's capabilities. The initial focus will be on the Environmental Control System (ECS), followed by expansion to include the Electrical Power System (EPS), Fuel System, and Engine System.
- Definition of Operational Scenarios and Fault Modes: For each selected subsystem, realistic operational scenarios will be defined, encompassing typical flight phases and environmental conditions. Furthermore, relevant fault modes for each subsystem will be identified, based on industry data, failure mode and effects analyses (FMEAs), and expert knowledge. These fault modes will represent a range of common and critical failures that can occur in aircraft systems.
- Specification of Performance Objectives and Constraints: For each use case, specific performance objectives for the sensor network will be defined. These objectives will be derived from the overall aim of enhancing aircraft health management and reducing MRO costs. They will include metrics such as diagnostic accuracy, fault detection time, and prognostic horizon. Concurrently, practical constraints, such as sensor cost limitations, weight restrictions, and power consumption budgets, will be specified to ensure the framework addresses real-world implementation challenges.
- Conceptual Model Development: For each use case, a conceptual model of the subsystem and its sensor network will be developed. This model outlines the key components of the subsystem, potential sensor locations,

types of sensors to be considered, and the data flow within the system. These conceptual models will serve as the basis for the implementation and testing phases.

The use case conceptualisation phase ensures that the MOSOF framework is developed and evaluated within a realistic and relevant context, enhancing the practical value and applicability of the research outcomes to the aerospace industry. The conceptualised use cases will be further elaborated upon in subsequent chapters, particularly in Chapters 3, 4, and 5, as they guide the framework implementation and evaluation.

1.4.3 Phase 3: Framework (MOSOF) Development

Phase 3 is the core of this research, focusing on the development of the Multi-Objective Sensor Optimisation Framework (MOSOF). Building upon the insights gained from the literature review (Phase 1) and the use case conceptualisation (Phase 2), this phase involves the following key activities:

- Formulation of Multi-Objective Cost Functions: This involves the mathematical formulation of cost functions that encapsulate the multiple, often conflicting, objectives of sensor optimisation. These cost functions will quantitatively represent performance criteria, including diagnostic accuracy, sensor network cost, weight, energy consumption, and coverage. The cost functions will be designed to be flexible and adaptable, allowing for customisation based on specific application requirements and priorities. The development of these cost functions will draw upon established methodologies in multi-objective optimisation and system evaluation.
- Development of the Sensor Optimisation Algorithm: This step will focus on selecting and/or developing an appropriate optimisation algorithm to effectively search the sensor configuration space and identify Pareto-optimal solutions that balance the defined cost functions. The algorithm will be chosen based on its suitability for multi-objective optimisation problems, its computational efficiency, and its ability to handle the complexities of sensor network design. Potential algorithms to be considered include evolutionary algorithms, gradient-based methods, and hybrid approaches.
- Integration of the Normalised Diagnostic Contribution Index (NDCI): A key feature of the MOSOF framework will be the integration of a Normalised Diagnostic Contribution Index (NDCI). The NDCI will be specifically designed to quantify the diagnostic contribution of individual sensors or sensor combinations to the overall system health assessment. This index will be formulated to be sensitive to various fault conditions and to provide a normalised measure of diagnostic information gain. The NDCI will be integrated into the MOSOF framework to guide sensor selection and placement decisions, ensuring that the optimised sensor network effectively contributes to accurate fault diagnosis.

- Framework Architecture and Implementation: This activity involve defining the overall architecture of the MOSOF framework and developing its computational implementation. The framework will be designed to be modular and extensible, allowing for the integration of different optimisation algorithms, cost functions, and diagnostic indices. The implementation will be carried out using appropriate software tools and programming languages, facilitating its application to the defined use cases.

The outcome of Phase 3 is a fully developed and computationally implemented MOSOF framework, ready for application and testing in the subsequent phases. The detailed design and formulation of the MOSOF framework, including the multi-objective cost functions and the NDCI, will be presented in Chapter 3, "Normalised Diagnostic Contribution Index Integration to Multi Objective Sensor Optimisation Framework (MOSOF) – An Environmental Control System Case."

1.4.4 Phase 4: Implementation and Testing

Phase 4 focuses on the practical implementation and testing of the MOSOF framework and the NDCI within the context of the defined aircraft subsystem use cases. This phase involves:

1. Simulation Environment: A simulation environment will be used to represent the selected aircraft subsystems (ECS, EPS, Fuel System, Engine System) and their sensor networks. This simulation environment will incorporate models of subsystem behaviour, sensor characteristics, and potential fault modes. The simulation platform will be chosen to accurately represent aircraft system dynamics and enable the generation of realistic sensor data.
2. NDCI-MOSOF Application to Multiple Aircraft Subsystems: Following the ECS application, the NDCI-MOSOF framework will be extended and applied to the other selected aircraft subsystems (EPS, Fuel System, Engine System). This multi-subsystem application will demonstrate the framework's versatility and scalability in handling different types of aircraft systems and diagnostic challenges. The implementation will involve adapting the cost functions and NDCI as necessary to suit the specific characteristics of each subsystem.
3. Performance Testing and Data Generation: Extensive simulation experiments will be conducted to test the performance of the implemented NDCI-MOSOF framework in each use case. These experiments will involve simulating various fault conditions, generating sensor data, and evaluating the diagnostic performance of the framework using relevant metrics, such as fault detection accuracy, false alarm rate, and diagnostic resolution. The generated data will be used for subsequent evaluation and analysis.

The implementation and testing phase provides empirical evidence of the MOSOF framework's practical applicability and performance in realistic aircraft system scenarios. The detailed implementation and testing procedures, along with the results obtained for the ECS and multiple aircraft subsystems, is presented in Chapter 4, "MOSOF with NDCI – A Cross-Subsystem Evaluation of an Aircraft for an Airline Case Scenario ."

1.4.5 Phase 5: Evaluation and Conclusion

The final phase of the research is dedicated to a comprehensive evaluation of the MOSOF framework and the NDCI, drawing conclusions based on the findings, and outlining directions for future work. This phase encompasses:

1. Performance Evaluation and Analysis: The performance data generated in Phase 4 will be rigorously analysed to evaluate the effectiveness of the MOSOF framework and the NDCI in achieving the defined objectives. This analysis will focus on assessing the framework's diagnostic accuracy, robustness, computational efficiency, and ability to balance multi-objective performance criteria. The evaluation will utilise appropriate statistical methods and performance metrics to quantify the framework's strengths and limitations.
2. Use Case Validation: The results obtained from the use case scenarios will be used to validate the practical applicability and relevance of the MOSOF framework to real-world aircraft health management challenges. The validation will assess the framework's ability to address the identified gaps in current approaches and its potential to contribute to improved aircraft maintenance practices.
3. Framework Limitations and Strengths Assessment: Based on the evaluation and validation results, a critical assessment of the MOSOF framework's limitations and strengths will be conducted. This assessment will identify areas where the framework performs effectively, as well as areas requiring further improvement or refinement. The limitations and strengths will be discussed in the context of the research objectives and the broader field of sensor optimisation and complex systems health management.
4. Conclusion and Future Work Directions: The research will be concluded by summarising the key contributions, findings, and insights gained throughout the study. Based on the evaluation and limitations assessment, specific directions for future research and development will be outlined. These future work directions will aim to address the identified limitations, extend the framework's capabilities, and explore new application areas, building upon the foundations established by this thesis.

The outcomes of Phase 5 will be synthesised and presented in Chapter 5, "Conclusion, Contributions and Future Work." This chapter provides a comprehensive evaluation of the research, draws definitive conclusions, and charts a course for future advancements in the field.

1.5 Structure of the thesis

This thesis is structured to provide a logical and progressive exposition of the research, commencing with foundational knowledge and culminating in the presentation of findings, conclusions, and future directions. The thesis is organised into the following chapters:

Chapter 2: Understanding the Role of Sensor Optimisation in Complex Systems [31]

This chapter lays the theoretical groundwork for the entire thesis by presenting a comprehensive literature review. It delves into the fundamental principles of sensor optimisation within the context of complex systems, exploring the state-of-the-art methodologies, frameworks, and techniques currently employed. The chapter examines multi-objective sensor optimisation approaches, diagnostic indices used for system health assessment, and the application of sensor optimisation in aerospace and similar complex domains. Critically, this chapter identifies existing research gaps and establishes the specific niche that this thesis aims to address. It provides the necessary context and theoretical underpinning for the development of the Multi-Objective Sensor Optimisation Framework (MOSOF) presented in subsequent chapters.

Chapter 3: Integrating a Normalised Diagnostic Contribution Index within the MOSOF Framework for Environmental Control Systems (ECS) [32]

Building upon the foundational understanding established in Chapter 2, Chapter 3 details the design and formulation of the Multi-Objective Sensor Optimisation Framework (MOSOF). It elaborates on the principles of system-aware sensor design and presents the development of multi-objective cost functions that are integral to the framework. The chapter specifically focuses on the integration of a novel Normalised Diagnostic Contribution Index (NDCI) within MOSOF. It then demonstrates the initial application of the MOSOF framework, focusing on the development of a diagnostic index tailored explicitly for the health assessment of aircraft Environmental Control Systems (ECS). This chapter provides a detailed explanation of the framework's architecture, its key components, and the mathematical formulations underpinning its functionality in the context of ECS diagnostics.

Chapter 4: Applying the NDCI-MOSOF Framework to Enhance Diagnostics Across Multiple Aircraft Subsystems [33]

Chapter 4 expands the scope of the MOSOF framework beyond its initial application in ECS. It details the application of the NDCI-MOSOF framework to enhance diagnostic capabilities across a broader range of critical aircraft subsystems. Specifically, this chapter presents the implementation and testing of the framework for the Electrical Power System (EPS), Fuel System, Engine System, and the Environmental Control System, now considered within a vehicle-level diagnostic perspective. This chapter showcases the framework's versatility and scalability in handling different types of aircraft systems and diagnostic challenges. It presents the results of performance testing within simulation environments, demonstrating the framework's effectiveness in a multi-subsystem context.

Chapter 5: Conclusion, Contributions, and Directions for Future Research

Concluding the thesis, Chapter 5 provides a comprehensive summary of the research undertaken, highlighting the key contributions to the field of sensor optimisation and complex systems health management. It revisits the study's aim and objectives, outlining the extent to which they have been achieved. The chapter critically examines the limitations of current research and proposes specific, actionable directions for future work. These future research directions aim to build upon the foundations established by this thesis, further refine the MOSOF framework, and explore new avenues for its application and development in the evolving landscape of aerospace health management and beyond.

This structured approach ensures a clear and coherent narrative, guiding the reader through the research process from motivation and theoretical foundations to framework development, application, validation, and ultimately, to the conclusions and future prospects of this work.

1.6 List of Publications

Journal Papers:

1. Suslu, B.; Ali, F.; Jennions, I.K. Understanding the Role of Sensor Optimisation in Complex Systems. Sensors 2023, 23, 7819.
2. Suslu, B.; Ali, F.; Jennions, I.K. Normalised Diagnostic Contribution Index (NDCI) Integration to Multi Objective Sensor Optimisation Framework (MOSOF)—An Environmental Control System Case. Sensors 2024, 24, 2661.
3. Suslu, B.; Ali, F.; Jennions, I.K. MOSOF with NDCI: A Cross-Subsystem Evaluation of an Aircraft for an Airline Case Scenario. Sensors 2026, 26, 160.

1.7 References

1. International Air Transport Association (IATA), "Airline profitability to strengthen slightly in 2025 despite supply chain challenges," *Press Release*, Jun. 2, 2025. https://www.iata.org/en/pressroom/2025-releases/2025-06-02-01
2. IATA, "Strengthened profitability expected in 2025 even as supply chain challenges persist," *Press Release*, Dec. 10, 2024. https://www.iata.org/en/pressroom/2024-releases/2024-12-10-01/
3. IATA, "Airline profitability outlook improves for 2024," *Press Release*, Jun. 3, 2024. https://www.iata.org/en/pressroom/2024-releases/2024-06-03-01/
4. Air Transport Action Group (ATAG), "Facts & figures (global aviation overview)," 2025. https://atag.org/facts-figures
5. Oliver Wyman, *Global Fleet and MRO Market Forecast 2024–2034*, 2024. https://www.oliverwyman.com/our-expertise/insights/2024/feb/global-fleet-and-mro-market-forecast-2024-2034.html
6. Aeronautical Repair Station Association (ARSA) / Oliver Wyman, *Global Fleet & MRO Market Report—Executive Summary*, 2024. https://arsa.org/wp-content/uploads/2024/03/ARSA-OW-2024MarketReport-ExecutiveSummary.pdf
7. FAA/NEXTOR (M. Ball, et al.), *Total Delay Impact Study: A Comprehensive Assessment of the Costs and Impacts of Flight Delay in the United States*, 2010. https://isr.umd.edu/NEXTOR/pubs/TDI_Report_Final_10_18_10_V3.pdf
8. EUROCONTROL, *Standard Inputs for Benefit-Cost Analyses in the ATM Context*, Ed. 9.0, 2024. https://www.eurocontrol.int/sites/default/files/publication/files/standard-input-for-eurocontrol-cost-benefit-analyses-2018-edition-8-version-2.6.pdf
9. Airbus, "Artificial intelligence for the aerospace sector" (Airbus A350 sensor and data figures), 2016.

https://www.airbus.com/en/newsroom/news/2016-12-artificial-intelligence-for-the-aerospace-sector

10. Royal Aeronautical Society, "The digital takeover of aircraft" (flight-control parameterisation overview), 2018. https://www.aerosociety.com/news/digital-takeover/
11. I.K. Jennions (Ed.), *Integrated Vehicle Health Management: Perspectives on an Emerging Field (R-405)*. Warrendale, PA: SAE International, 2011. https://www.sae.org/books/integrated-vehicle-health-management-perspectives-emerging-field-r-405/
12. I.K. Jennions (Ed.), *Integrated Vehicle Health Management: The Technology (R-429)*. Warrendale, PA: SAE International, 2013. https://www.sae.org/books/integrated-vehicle-health-management-technology-r-429/
13. C. M. Ezhilarasu, Z. Skaf, and I. K. Jennions, "The application of reasoning to aerospace IVHM: Challenges and opportunities," *Progress in Aerospace Sciences*, vol. 105, pp. 60–106, 2019, https://doi.org/10.1016/j.paerosci.2019.01.001
14. L. Basora, X. Olive, and T. Dubot, "Recent advances in anomaly detection methods applied to aviation," *Aerospace*, vol. 6, no. 11, p. 117, 2019, https://doi.org/10.3390/aerospace6110117
15. M. Bieber, W. J. C. Verhagen, and B. F. Santos, "A generic diagnostic framework for anomaly detection," *Aerospace*, vol. 10, no. 8, p. 673, 2023, https://doi.org/10.3390/aerospace10080673
16. M. J. Page *et al.*, "The PRISMA 2020 statement: An updated guideline for reporting systematic reviews," *BMJ*, vol. 372, n71, 2021, https://doi.org/10.1136/bmj.n71
17. H. Peng, F. Long, and C. Ding, "Feature selection based on mutual information: Criteria of max-dependency, max-relevance, and min-redundancy," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 27, no. 8, pp. 1226–1238, 2005, doi:10.1109/TPAMI.2005.159
18. J. R. Vergara and P. A. Estévez, "A review of feature selection methods based on mutual information," *Neural Computing and Applications*,

vol. 24, no. 1, pp. 175–186, 2014, https://doi.org/10.1007/s00521-013-1368-0

19. M. S. Daigle, I. Roychoudhury, and A. Bregon, “A diagnosability-based approach to sensor selection using structural model decomposition,” in *Proc. IEEE Aerospace Conf.*, 2014, pp. 1–10, https://doi.org/10.36001/phme.2014.v2i1.1545
20. B. Chen, S. L. Shah, A. K. Jain, and T. Chen, “Effectivity analysis of adding sensors for improving fault isolability properties,” *Journal of Process Control*, vol. 69, pp. 134–146, 2018, https://doi.org/10.1016/j.jprocont.2018.07.007
21. K. Deb, *Multi-Objective Optimization Using Evolutionary Algorithms,* https://www.egr.msu.edu/~kdeb/papers/k2011003.pdf
22. K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, “A fast and elitist multiobjective genetic algorithm: NSGA-II,” *IEEE Trans. Evol. Comput.*, vol. 6, no. 2, pp. 182–197, 2002, https://dl.acm.org/doi/10.1109/4235.996017
23. International Organization for Standardization, *ISO 13374-4:2015—Condition monitoring and diagnostics of machine systems—Data processing, communication and presentation—Part 4: Presentation*, 2015, https://www.iso.org/standard/54933.html
24. MIMOSA, “Specifications listing” (OSA-CBM / OSA-EAI open standards for asset management and CBM), 2021–2025. https://www.mimosa.org/specifications-listing/
25. Object Management Group (OMG), “Open System Architecture for Condition-Based Monitoring (OSA-CBM) overview,” 2007, https://www.omg.org/mda/mda_files/OSACBM.htm
26. J. Coble and J. W. Hines, “Identifying optimal prognostic parameters from data: A genetic algorithms approach,” *Proceedings of the Annual Conference of the PHM Society*, 2009.
27. Y. Lei, N. Li, L. Guo, N. Li, T. Yan, and J. Lin, “Machinery health prognostics: A systematic review,”

Mechanical Systems and Signal Processing, vol. 104, pp. 799–834, 2018.

28. A. K. S. Jardine, D. Lin, and D. Banjevic, “A review on machinery diagnostics and prognostics implementing condition-based maintenance,” *Mechanical Systems and Signal Processing*, vol. 20, no. 7, pp. 1483–1510, 2006.
29. R. Zhao, R. Yan, Z. Chen, K. Mao, P. Wang, and R. X. Gao, “Deep learning and its applications to machine health monitoring,” *Mechanical Systems and Signal Processing*, vol. 115, pp. 213–237, 2019.
30. T. Wang, J. Yu, D. Siegel, and J. Lee, “A similarity-based prognostics approach for remaining useful life estimation,” *IEEE Transactions on Reliability*, vol. 57, no. 4, pp. 628–640, 2008.
31. B. Suslu, F. Ali, and I. K. Jennions, “Understanding the role of sensor optimisation in complex systems,” Sensors, vol. 23, no. 18, p. 7819, 2023.
32. B. Suslu, F. Ali, and I. K. Jennions, “Normalised Diagnostic Contribution Index (NDCI) integration to MOSOF—An Environmental Control System case,” Sensors, vol. 24, p. 2661, 2024.
33. B. Suslu, F. Ali, I. K. Jennions, MOSOF with NDCI: A Cross-Subsystem Evaluation of an Aircraft for an Airline Case Scenario. *Sensors* 2026, *26*, 160.

2 Understanding the Role of Sensor Optimisation in Complex Systems

Chapter 2 (based on the first research paper) provides a comprehensive literature review on sensor optimisation in complex systems, with an emphasis on the context of Integrated Vehicle Health Management (IVHM). It surveys state-of-the-art techniques for determining an optimal set of sensors to monitor and diagnose complex assets (like aircraft) under multiple objectives. The review outlines how sensor selection in IVHM involves balancing conflicting factors – maximising fault detection and isolation capability while minimising costs, weight, and complexity. Traditional approaches and metrics (e.g., information gain methods, mutual information, redundancy–relevance algorithms like mRMR) are examined alongside newer multi-objective optimisation and decision-making techniques. The chapter introduces a taxonomy of sensor optimisation methods and discusses key performance measures (such as fault detection rate, isolation rate, false alarm rate) used to evaluate sensor sets. Crucially, Chapter 2 identifies a gap in the literature. While numerous studies address sensor placement and selection at the component or subsystem level, very few tackle the system-of-systems (platform-level) optimisation for health monitoring. In other words, past research has seldom considered interactions across multiple subsystems within an integrated platform. This insight highlights the need for holistic sensor optimisation strategies that account for cross-subsystem fault effects. The chapter's findings and identified gaps set the stage for the subsequent research, underscoring the potential benefits of a multi-subsystem approach to sensor suite design for improved safety, reliability, and maintenance efficiency in complex engineering systems.

2.1 Introduction

2.1.1 Background to the Literature

Health management applications are critical for ensuring the reliable and effective operation of mission and safety-critical complex engineering systems, such as spacecraft, submarines, aircraft, and industrial plants. These applications rely on sensors that generate useful information about the health condition of the assets, thus optimising the sensor network quality while considering specific constraints. This is the first step in assessing the condition of assets. This is particularly important for systems that operate in harsh environments, where the accuracy and reliability of sensor data are essential for effective maintenance and operation.

Integrated Vehicle Health Management (IVHM) involves monitoring, assessing, and predicting the health of various systems within a vehicle or a larger system. This approach aims to provide a comprehensive understanding of the system's health by integrating data from multiple sensors and other sources and analysing it using advanced algorithms and analytical techniques to support Condition-Based Maintenance (CBM) [1]. In recent years, there has been a growing interest in sensor optimisation techniques for IVHM, which involve optimising sensor placement, selection, and data processing to enhance the accuracy and reliability of the health monitoring system.

Sensing technology is an essential tool for understanding and interacting with the physical world. By providing accurate and timely data, sensors allow the algorithms to make decisions and reason about the physical phenomena in an environment. IVHM systems build on this to provide accurate and comprehensive information about the system's state, utilising an optimal sensor suite that contributes to increased system availability, effective reasoning for performance and maintenance actions. The sensor optimisation problem in complex systems involves finding the most effective combination of sensor types and locations that will provide the most accurate and comprehensive monitoring of a system.

To optimise sensor selection for IVHM or complex systems, various factors must be considered, including the cost, weight, size, and power consumption of the sensors. These factors can impact the feasibility and practicality of using specific types of sensors, potentially affecting the overall performance and accuracy of the IVHM system. Historically, these objectives, from an engineering perspective, are viewed as the sensor optimisation problem; however, this study takes a broader view, including the information gain from the sensor suite.

Additionally, the optimisation problem in sensor selection also involves considering trade-offs between different performance metrics. For example, from information theory, increasing the number of sensors may improve the accuracy of the IVHM system. Still, it may also increase the weight, cost and computational complexity of the system. Balancing these trade-offs requires careful consideration of the specific requirements and constraints of the application, as well as a deep understanding of the underlying physics and dynamics of the system being monitored. Another essential factor in sensor selection is the type of data that is required for fault detection and diagnosis. Different types of sensors may be better suited for monitoring various aspects of a vehicle's health, including temperature, pressure, vibration, and fluid levels. By selecting the right combination of sensors, IVHM systems can ensure that they detect potential faults accurately and provide timely alerts to maintenance personnel.

This review intends to provide a comprehensive overview of the existing literature on sensor optimisation for complex systems. It will explore various optimisation techniques used in different domains and highlight the key findings and limitations of each approach. Additionally, the review will identify areas for future research and discuss potential benefits of sensor optimisation for IVHM systems.

2.1.2 Problem Statement

The use of IVHM systems in aircraft operation has become increasingly important for ensuring safety, increasing operational efficiency, and reducing downtime and maintenance costs. However, the effectiveness of these systems relies on identifying optimal sensors and their optimal locations. Sensors are used to translate physical phenomena into digital form, and the optimisation of these sensors can be achieved in three different ways: physical, system, and algorithmic. The first approach involves improving the quality of the sensing technology itself. The second approach combines location and quantity to achieve better information quality. The last approach involves enhancing processing techniques.

The number of heterogeneous parameters in flight data, collected from different types of sensors in the aircraft, is increasing due to high safety requirements, incident and accident investigation, maintenance, and diagnostic purposes. Analysing the entire sensor data, up to thousands of parameters in modern aircraft, is neither practical nor computationally manageable for onboard diagnostic purposes. The data types range from Boolean data to control systems' BITE codes, as well as high-frequency vibration or acoustic data. The analysis of the latter, which relies on defining which features in the data have more importance in terms of clear condition indicators, is an open research area.

The optimal sensor suite identification problem in IVHM systems involves determining the sensor suite that will provide the most comprehensive and accurate monitoring of the aircraft`s health situation. This requires consideration of various factors, including the type of data that is necessary for fault detection and diagnosis, the cost and weight of sensors, and power consumption. While significant research has been conducted in this area, further investigation is needed in the system-level optimisation of sensor suites for IVHM systems.

There are numerous research papers published on aircraft component and subsystem-level sensor optimisation, yet only a few papers consider system-level condition monitoring in aircraft. There is a need for further investigation into the use of advanced optimisation techniques and the integration of multiple sub-systems for more comprehensive health management. Different complex engineering systems that utilise sensor optimisation techniques will be investigated to improve the current research on the IVHM domain.

Additionally, other complex system applications that benefit significantly from sensor optimisation include wind turbines, power generation plants, railways, and satellites. Kulkarni et al. discussed and outlined a framework for sensor selection framework which was particularly designed for diagnostic applications and used wind turbines as their experiment [2].

Figure 2-1 illustrates fundamental complex engineering systems that heavily rely upon high-precision sensors to maintain their operations in a healthy state through constant monitoring. As depicted, these macro-systems encompass large-scale, mission-critical infrastructures such as nuclear and thermal power generation plants, electrical distribution substations, offshore wind farms, and hydroelectric dams. The immense scale, high capital cost, and operational criticality of these facilities dictate that unexpected failures can lead to severe economic and safety consequences. Consequently, these environments serve as prime candidates for advanced sensor optimisation frameworks to enable robust condition-based maintenance.

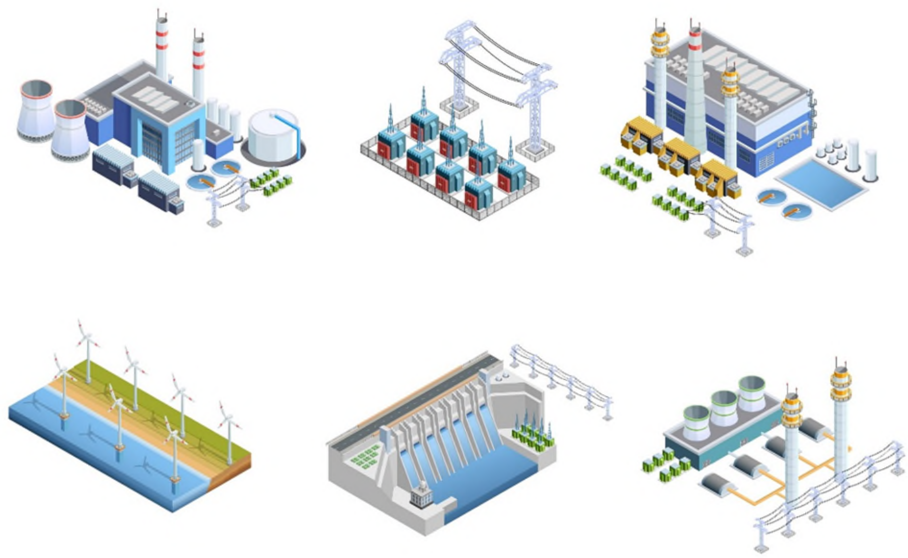

Figure 2-1 Some examples of the Complex Systems

Building upon this macro-level perspective, Figure 2-2 details the specific hierarchical flow of these complex systems, demonstrating their structural decomposition from the platform to the physical level. The diagram illustrates a top-down taxonomy, beginning with platform-level 'Complex Systems' (e.g., Aircrafts, Nuclear Power Stations, Wind Turbines, and Deep Space Satellites). These platforms are broken down into their functional 'Subsystems' (such as Engines, ECS, Power Units, and Control Units). The hierarchy then descends to the 'Components' level, which includes the physical hardware driving the system and ultimately terminates at the base 'Materials' level. From a diagnostic standpoint, this hierarchical mapping is critical; it visually reinforces the premise that optimal sensor selection cannot occur at the component level in isolation. Instead, diagnostic data generated by sensors must be evaluated based on how effectively it captures subsystem interactions and flows upward through the hierarchy to provide an accurate health assessment of the entire complex system.

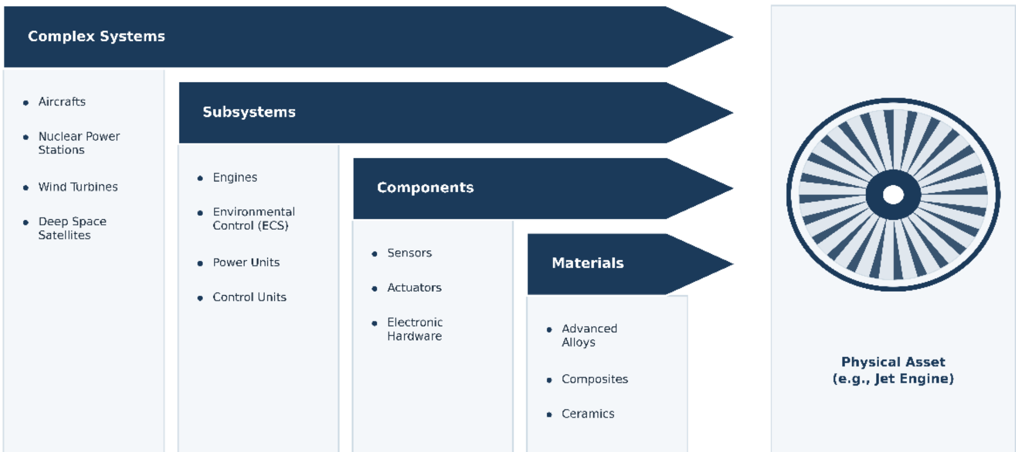

Figure 2-2 Direction and Examples of the Complex Systems` Themes

Defining a step-by-step methodology to rank/identify aircraft sensors for system-level condition monitoring and optimisation of the sensors for onboard/online diagnosis, where several domain-specific objective functions and constraints occur, is addressed in this study. Additionally, faulty sensor detection, which comprises most of the faults in the aviation industry, is mainly neglected in the optimal sensor suite selection research; this issue will also be considered to be integrated into the sensor suite selection part.

The motivation for further research in sensor optimisation for IVHM systems is driven by the potential benefits of more accurate and reliable on-board diagnosis for health monitoring systems. This could lead to improved safety and availability, as well as efficient CBM. Additionally, advances in sensor technologies and optimisation techniques present new opportunities for many different fields. The research will also identify areas for future investigation and highlight the potential benefits of a detailed sensor optimisation strategy in IVHM and provide guidance for researchers and practitioners in the field.

2.1.3 Scope of the Study

This research focuses on exploring the existing literature on sensor optimisation for placement, selection, operational applications, and data processing techniques used to improve diagnostic information.

A systematic and thorough approach to sensor selection, ensuring the data collected is relevant and valuable in the target application, must be carefully designed. The following questions need to be considered in the design phase:

What is the precise application for which the sensor data will be utilised? This will help guide the selection of suitable sensors and data collection methods.

What are the key variables or parameters that need to be measured to achieve the desired outcome? For example, if the goal is to monitor the performance of a machine, collecting data on factors such as temperature, vibration, and power consumption is needed.

What is the required level of accuracy for the data? This will help determine the appropriate resolution and precision for the sensors.

What is the frequency at which data needs to be collected? This will impact the selection of sensors and the required data storage and processing capabilities.

What are the environmental factors that may affect sensor performance, such as temperature, humidity, and electromagnetic interference? This will help guide the selection of sensors that can operate reliably under the specific conditions of the application.

By answering these fundamental questions and conducting a thorough analysis of the target system and its specific applications, a framework can be designed to ensure that the selected sensors are suitable for the task at hand and that the resulting data is accurate and helpful for analysis and decision-making.

2.2 Structuring the Review

2.2.1 Methodology of the Review

To conduct a literature review on sensor optimisation in the complex system domain, electronic databases, including scientific journals, conference proceedings, and relevant books, were searched by using keywords related to the topic. The articles and books were then screened for their relevance to the research gaps, and only those that met the inclusion criteria were selected for review.

The inclusion criteria for this review are studies that examine sensor optimisation in complex systems to improve the quality of diagnostic information, particularly those used in the aviation domain. The physical improvement of the sensing technology itself and the structural health management (SHM) field are beyond the scope of this research. Several SHM examples are included to demonstrate the use of the proposed method.

In this review, our methodology encompasses a systematic literature selection process, including database searches and well-defined inclusion criteria, to gather pertinent sensor optimisation research. The data from selected sources, enabling a comparative analysis of sensor optimisation techniques, application domains, and future trends are meticulously extracted and categorised. Our review methodology culminates in a concise summary of key findings, emphasising the practical implications of sensor optimisation across diverse domains. This structured approach ensures methodological rigour and contributes to advancing the understanding of sensor optimisation techniques within complex systems.

Table 2-1 shows search results from the Scopus data. Search terms included various combinations of keywords.

Table 2-1 Search Results in Scopus Database

TITLE – ABSTRACT- KEYWORDS	Document Results
(Sensor AND Optimisation)	75.169
(Sensor AND Optimisation) AND (Complex AND Systems)	10.103
(Sensor AND Optimisation) AND (Complex AND Systems) AND (aircraft)	602
(Sensor AND Optimisation) AND (Complex AND Systems) AND (aircraft) AND (diagnostic)	85
“Sensor Optimisation”	509
“Sensor Optimisation” AND “Complex Systems” OR “aircraft”	63
“Sensor Optimisation” AND “Complex Systems”	9

Without any restrictions, a search on sensor and optimisation yielded 75169 articles which are based on their titles, abstracts and keywords. When the other identified keywords of the study were added to this number dropped, as shown in the table, until a result with only 85 articles was found. When a stricter grouping was explored (by using quotation marks), it resulted in 9 articles that consider different complex systems.

After analysing the relevance and contribution of each article to the field, the information found is synthesised in order to categorise the articles into different themes and analyse the main findings and trends. Identified associated paper references were also searched to expand the knowledge that was gathered over the field, and additional studies that met the inclusion criteria are included.

2.2.2 Taxonomy

A taxonomy that was developed through a comprehensive analysis of the results obtained from a Scopus search, as well as relevant books and papers, is presented in Table 2-1. The primary objective of this taxonomy is to classify and organise the various approaches used to optimise sensor suites for complex systems. The taxonomy's classification criteria are based on the key domains and techniques employed in the process of sensor suite optimisation.

At the highest level of categorisation, the taxonomy consists of several primary branches, each representing a significant area or approach used in sensor optimisation. These branches serve as the fundamental divisions, providing an overarching structure to the taxonomy. Subsequently, within each main branch, further divisions and subcategories are established by posing fundamental investigative questions. These questions are designed to explore and dissect the various aspects and intricacies of optimising sensor suites for complex systems comprehensively.

By employing this matrix organisation table in Figure 2-3, the taxonomy facilitates a systematic and methodical examination of the field, enabling researchers and practitioners to navigate through the vast array of sensor optimisation methods and concepts effectively.

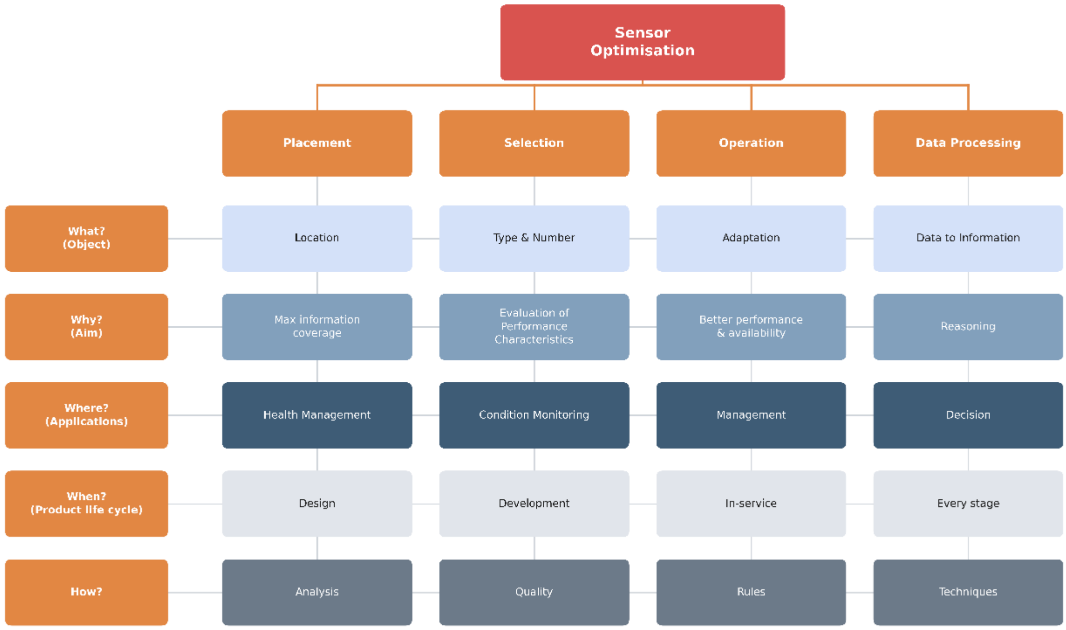

Figure 2-3 Taxonomy of Sensor Optimisation

This approach not only provides a comprehensive overview of the field but also offers actionable insights and recommendations for enhancing sensor performance and capabilities within complex systems. The initial taxonomy serves as an inclusive framework, laying the groundwork for subsequent sections to delve into sensor improvement areas. By leveraging the strengths of approaches, the paper can deliver a well-rounded exploration of sensor optimisation while offering targeted recommendations for improvement wherever appropriate.

Figure 2-3 delineates the chronological lifecycle of a sensor network's implementation, advancing sequentially from initial placement and selection through to operational deployment and downstream data processing. Building upon this framework, Figure 2-4 expands specifically on these aspects' sub sections to create the conceptualisation that govern the latter stages of this lifecycle. It is important to note that while Figure 2-3 maps the temporal progression of physical engineering deployment, the subsequent literature review intentionally adopts a thematic, rather than chronological, structure. This thematic arrangement serves to isolate and prioritize the fundamental architectural bottlenecks of the design process. Consequently, the review focuses first on the complex multi-objective optimisation algorithms requisite for sensor selection and placement, before systematically addressing the subsequent data processing and diagnostic methodologies.

Figure 2-4 presents a concept map that has been thoughtfully organised to align with the key themes derived from the comprehensive literature review. Each major title in the concept map corresponds to a significant area of investigation, facilitating a clear understanding of the interrelationships and connections between the various concepts and topics explored in the reviewed literature.

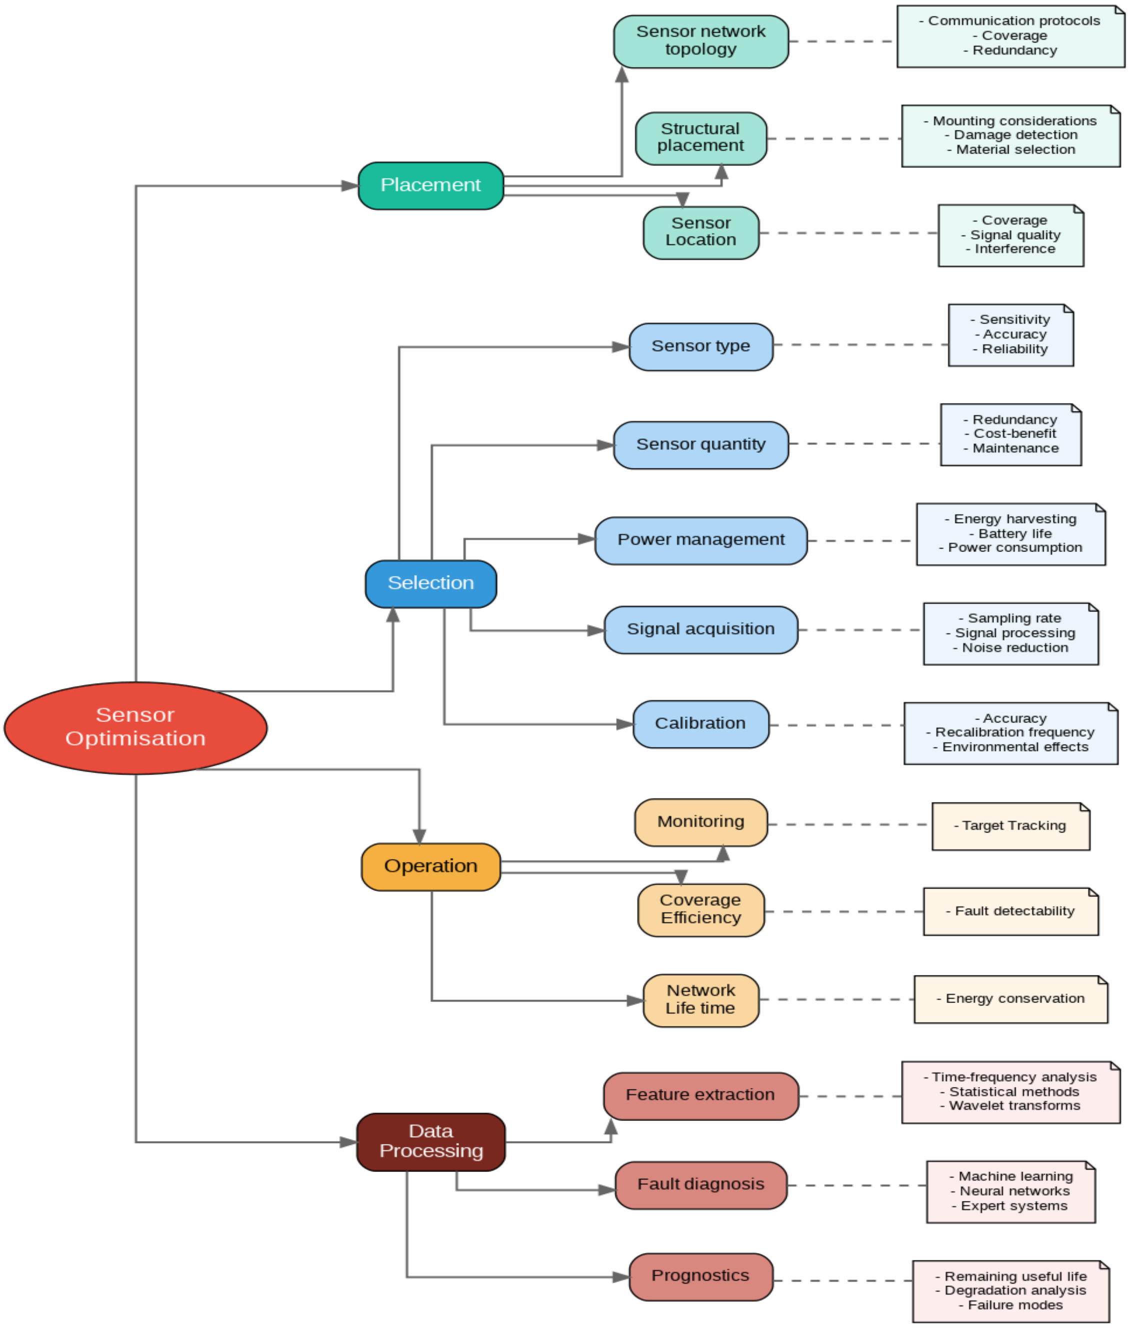

Figure 2-4 Conceptual Map of the Literature

The concept map visually organises the main themes from the literature review, providing a coherent and interconnected representation of the key concepts and considerations in sensor optimisation within complex systems. These brief descriptions offer readers a clear roadmap for further exploration and understanding of the topics presented in Figure 2-4. Short descriptions of the main branches were then given in the following paragraphs.

Placement: The placement of sensors requires consideration of both their physical location and how they are mounted or integrated into the system. The fundamental goal of sensor positioning is to make sure that sensors are ideally situated to gather the necessary data to assess the system's health information. This category is further divided into structural placement, network topology, and sensor location subcategories. According to the system's structural characteristics, the first entails deciding how the sensors should be placed, the second deals with how the sensors should be set up in a network to ensure total coverage and redundancy, while the last one deals with location and interference.

Selection: The type and number of sensors employed in the system must all be considered throughout the selection process. Ensuring that the chosen sensors are suitable for the task and deliver accurate and reliable data is the primary goal of sensor selection. To this end, several factors need to be assessed accordingly, including power management, signal acquisition, and calibration. When choosing a specific sensor type, consideration must be given to sensitivity, accuracy, and dependability. Sensor quantity entails selecting the correct number of sensors to balance redundancy, information gain, and cost.

Operation: Considering how sensors are powered, how data is acquired, and how they are calibrated to preserve accuracy over time are all part of the operation. The fundamental goal of the operation is to make sure that sensors are accurate and performing well. There are three subcategories within this category: network lifetime, coverage efficiency, and monitoring. To ensure reliable operations, coverage efficiency over the network lifetime of the sensors should be managed precisely. It is imperative to capture the necessary information and adjust system parameters in response to changing conditions to inform decision support systems through continuous monitoring.

Data Processing: To identify errors and forecast system behaviour, data processing considers factors related to how sensor data is processed, examined, and interpreted. The primary goal of data processing is to convert sensor data into useful information that can be applied to informed decision-making. With sound data processing, the best use of the data can be made with appropriate techniques to maximise the information gain. The three subcategories within this area are feature extraction, fault diagnosis and prognostics. To extract pertinent characteristics from sensor data, feature extraction uses time-frequency analysis, statistical techniques, and wavelet transforms. Machine learning, neural networks, and expert systems are utilised in fault diagnostics to identify systematic patterns. Prognostics includes estimating the system's remaining useful lifetime, examining degradation, and determining failure modes.

Overall, the proposed taxonomy provides a comprehensive framework for understanding the key concepts and considerations related to sensor optimisation in IVHM systems. This is also useful for organising and synthesising the literature, as well as for identifying patterns and trends in the research domain and providing guidance for future research directions. The remainder of this paper presents a detailed review of the four major areas identified, accompanied by specified cost function propositions for each aspect, and concludes with a discussion on future research directions.

2.3 Placement

Optimal sensor placement (OSP) aims to determine the optimal number and location to be deployed in a system, while considering various factors such as cost, measurement accuracy, and system performance. The use of OSP is particularly important in systems where the placement of sensors can significantly impact the diagnostic accuracy and performance of the system being monitored. For instance, in fault diagnosis systems, improper sensor placement can lead to inaccurate fault detection and diagnosis.

2.3.1 Theoretical Background and OSP Methods

Various theoretical principles are used in OSP, including statistical analysis, mathematical modelling, and simulation techniques. One key aspect of OSP is understanding the key parameters that affect optimal sensor placement and the trade-offs between objectives and constraints. Another important aspect is the use of inverse problems to determine the location of a signal source from sensor measurements. Modelling sensor data using Gaussian Process (GP) is also a commonly used technique in OSP, which is a natural generalisation of linear regression that allows for the consideration of uncertainty about predictions. The comparison of the main techniques used in OSP is shown in Table 2-2.

Table 2-2 Comparison of Fundamental OSP Techniques

Optimisation Technique	Pros	Cons
Statistical Analysis	Provides valuable insights from data analysis	May require large datasets and complex statistical methods
Mathematical Modelling	Enables accurate representation of system behaviour	Can be computationally intensive for complex systems
Simulation Techniques	Allows testing in virtual environments	May not fully capture real-world complexities
Inverse Problems	Determines source location from sensor measurements	Can be sensitive to measurement errors and noise
Gaussian Process (GP)	Accounts for uncertainty in predictions	May require significant computational resources

Nakai *et al*. [3] discussed different approaches to sensor placement optimisation for various applications. The focus was on selecting the optimal set of sensors to estimate high-dimensional data based on different optimality criteria, such as D-, A-, and E-optimality, which are used to maximise the determinant, minimise the trace of the inverse, and maximise the minimum eigenvalue of the Fisher information matrix, respectively. The performance of the greedy algorithms based on these criteria was evaluated using randomly generated systems and a practical dataset related to climate science. The comparison of the two approaches` pros and cons is stated in Table 2-3.

Table 2-3 Comparison of the Optimality Criteria and Greedy Search Algorithms

Approach	Pros	Cons
D-, A-, E-optimality	Evaluates optimality based on specific criteria	May not consider other important aspects of sensor placement
Greedy Algorithms	Relatively simple to implement	May not always find the global optimum

Wan *et al.* [4] reviewed the optimal sensor placement for Aircraft Structural Health Management (ASHM), which mainly focuses on structural health monitoring and assessment, microstructure fault monitoring and isolation, overload, corrosion monitoring, and residual life assessment. The authors emphasise the importance of OSP for ASHM and discuss the difficulty of optimising sensor placement in large aircraft structures. The study proposes the use of Singular Value Decomposition (SVD), QR decomposition, and Fuzzy Measurement Coverage to optimise sensor measuring points. The final result of the OSP is verified through QR decomposition and Fuzzy Measurement Coverage, and the OSP scheme is analysed for aircraft wing structures. The comparison of the approaches is shown in Table 2-4.

Table 2-4 Comparison of the Decomposition-Based OSP Methods

Approach	Pros	Cons
Singular Value Decomposition	Useful for large-scale structures	May not be applicable to all types of systems
QR Decomposition	Efficient for sensor placement	Limited to certain types of optimisation problems
Fuzzy Measurement Coverage	Incorporates uncertainties in coverage	Requires careful tuning of fuzzy parameters

Krause *et al.* [5] also discussed the problem of choosing sensor locations when monitoring spatial phenomena modelled as Gaussian processes (GPs). It notes several common strategies and tackles the combinatorial optimisation problem of maximising the mutual information between the chosen locations and the locations that are not selected. The paper proves that this problem is NP-complete but describes a polynomial-time approximation that is within (1−1/e) of the optimum by exploiting the sub-modularity of mutual information. The paper extends its algorithm to exploit lazy evaluations and local structure in the GP, yielding significant speedups. It also extends the approach to find placements that are robust against node failures and uncertainties in the model, again exploiting the sub-modularity of the objective function. Finally, the paper demonstrates the advantages of the approach for optimising mutual information in an extensive empirical study on two real-world datasets. The comparison of the techniques is shown in Table 2-5.

Table 2-5 Comparison of the Information-Theoretic OSP Methods

Approach	Pros	Cons
Mutual Information	Exploits sub-modularity for efficiency	NP-complete problem, may not find the global optimum
Lazy Evaluations	Speeds up the optimisation process	May not fully capture all system dynamics
Robust Placements	Considers uncertainties and node failures	Requires additional computational complexity

The reviewed studies emphasise the importance of OSP in improving fault diagnosis efficiency, sensor layout selection, and ASHM. Each study proposes a different approach to optimising sensor placement, including the use of a dynamic fault tree, a dynamic Bayesian network, SVD, QR decomposition, and Fuzzy Measurement Coverage. They demonstrate the complex nature of OSP and highlight the need for a systematic approach to sensor placement, considering different criteria and factors such as the number of sensors, edge effect, measurement degree of freedom, similarity of sensor locations, and hiddenness of fault position, and provide insights which can be applied to various fields such as aerospace, energy, and transportation.

OSP methods have been widely used in various fields to solve the sensor placement problem. These methods can be broadly categorised into three groups: heuristic, evolutionary and deterministic. The optimisation of sensor placement is a challenging task that involves finding the optimal locations for sensors based on a set of objectives and constraints. A range of optimisation techniques has been developed to address this problem, including sensitivity-based and topology-based approaches, as well as linear and non-linear optimisation methods.

Some of the most used optimisation techniques for sensor placement include evolutionary algorithms, particle swarm optimisation, and greedy algorithms. Evolutionary algorithms, such as genetic algorithms or differential evolution, are based on the principles of natural selection and survival of the fittest. These algorithms generate a population of candidate solutions and iteratively improve them by applying genetic operators, such as mutation, crossover, and selection. Particle swarm optimisation is a population-based optimisation technique inspired by the behaviour of bird flocks or fish schools. It involves iteratively adjusting the position and velocity of a set of particles to find the optimal solution. Greedy algorithms are simple heuristic techniques that aim to find the optimal solution by iteratively adding or removing sensors based on a set of criteria, such as the information gain or the cost-benefit ratio.

Sensitivity-based approaches aim to identify the most sensitive locations in a system by analysing the system's response to changes in input parameters. This involves calculating the sensitivity coefficients, which describe how changes in the sensor measurements affect the system output. These coefficients are used to guide the placement of sensors to maximise the system's response to changes in the input parameters.

Topology-based approaches focus on identifying the most critical locations in a system by analysing the system's topology. This involves identifying the most critical nodes, edges, or regions in the system using graph-based techniques, such as centrality analysis or clustering. These techniques can be used to guide the placement of sensors to maximise the coverage of the critical regions in the system.

In [6] Clark *et al.* discussed the problem of optimal sensor placement under a cost constraint, which arises in various industrial and scientific applications. A well-established greedy algorithm for optimal sensor placement without cost constraints is extended to incorporate cost constraints, and its effectiveness is demonstrated on datasets related to facial recognition, climate science, and fluid mechanics. The paper emphasises that the cost-error landscape varies by application, and intuitive connections to underlying physics are observed.

Gomes *et al. [32] investigated the issue of identifying structural damage in large-scale structures, particularly* in aerospace applications. A metaheuristic algorithm called the "firefly algorithm" (FA) is used to identify structural damage by solving an inverse problem. The Fisher information matrix is used to optimise sensor placement, and the results demonstrate that optimised sensors contribute to improved identification of damage, especially in complex and large-scale structures. The proposed optimised damage identification process using FIM-FA has the potential to be extended to a wide range of structural health monitoring (SHM) applications in complex structures, where traditional non-destructive inspection methods may not be practical due to the structure's complexity and restricted access.

Feng *et al.* [7] reviewed the development of an OSP scheme to improve fault diagnosis efficiency, considering common cause failure. The study introduces a dynamic fault tree converted to a dynamic Bayesian network to calculate reliability parameters and constructs the decision matrix. An efficient TOPSIS algorithm is adopted to determine the potential sensor locations. A diagnostic sensor model is also developed to account for the failure sequence between a sensor and a component. The authors provide a case study to prove the significant impact of common cause failure on sensor placement.

Yang *et al.* [8] focused on the investigation of OSP for a multirotary-joint solar power satellite (MJ-SPS) using six OSP methods to select the best sensor layout. Three standards and two novel criteria — sensor distribution and similarity of sensor locations — are added to evaluate the effectiveness of the sensor configurations. The study emphasises the importance of the work for the MJ-SPS and OSP methods, comparing different numbers of sensors and orders of modal shape.

It is important to note that the optimal sensor placement depends not only on the system's objectives but also on the type and number of sensors available, as well as the impact of environmental factors on sensor performance. For instance, the placement of temperature sensors in a heat exchanger may differ depending on the type of temperature sensor used (e.g., thermocouple, RTD, or infrared sensor) and the impact of fouling or corrosion on sensor accuracy.

Overall, the choice of sensor placement optimisation method depends on the complexity of the problem, the computational resources available, and the required level of solution accuracy. Each method has its advantages and disadvantages, and the optimal method should be selected based on the specific requirements of the problem. The comparison of the methods is shown in Table 2-6.

Table 2-6 Comparison of Heuristic and Evolutionary OSP Methods

Methods	Pros	Cons
Dynamic Fault Tree	Captures time-varying dependencies	Complexity increases with system size
Firefly Algorithm (FA)	Efficient for large-scale structures	May require parameter tuning
Dynamic Bayesian Network	Accounts for time dependencies	Requires an accurate model representation
Genetic Algorithms	Global search capability	Convergence may be slow
Particle Swarm Optimisation (PSO)	Converges quickly	May get stuck in local optima
Sensitivity-Based Approaches	Identifies sensitive locations	Sensitive to measurement errors
Topology-Based Approaches	Identifies critical locations	Complexity increases with system size

2.3.2 Sparsity and Data-Driven Learning

In this section, the use of sparsity and data-driven learning techniques in OSP are discussed. Sparsity has been extensively used to improve the performance of inverse problems by reducing the number of unknowns and increasing the robustness of the solution. Sparsity is a well-known concept in data science and optimisation theory. This concept has been widely applied in various fields, including image and signal processing, machine learning, and optimisation [9], [10], [11] .

In sensor placement optimisation, sparsity can be used to identify essential measurement points in complex systems, resulting in significant cost savings by reducing the number of sensors required. Data-driven learning approaches, such as compressed sensing and dimensionality reduction, have been used to achieve sparsity in sensor placement [12]. These methods have been successfully applied to various applications, including structural health monitoring, power system health monitoring, and water distribution networks.

Compressed sensing is a mathematical technique that allows the recovery of a sparse signal from a small number of measurements. Compressed sensing has also been used in conjunction with other optimisation methods, such as convex optimisation, to improve sensor placement performance. Dimensionality reduction is another data-driven approach that can be used for sensor placement optimisation. Dimensionality reduction is the process of reducing the number of variables in a dataset while retaining the most critical information. This can be achieved using techniques such as principal component analysis (PCA), proper orthogonal decomposition (POD) or T-distributed stochastic neighbour embedding (t-SNE). It can be used to identify important measurement points in a system and reduce the number of sensors required.

When dealing with nonlinear systems, standard techniques for feature selection and sensor placement that rely on linearity assumptions or simple statistical models can result in costly oversensing without guaranteeing the recovery of desired information from the measurements. To this end, Otto et al. [13] discuss the importance of sensor placement and feature selection in solving inverse problems in nonlinear systems, highlighting the limitations of existing techniques that rely on linearity or simple statistical models. To overcome these limitations, the authors propose a novel data-driven approach based on secant vectors between data points for a general type of nonlinear inverse problem. The approach is used to develop three efficient greedy algorithms that provide different robust and near-minimal reconstruction guarantees. The algorithms are demonstrated on two issues where linear techniques fail: sensor placement for reconstructing a fluid flow with a complex shock-mixing layer interaction and selecting fundamental manifold learning coordinates on a torus.

Overall, the application of sparsity and data-driven learning techniques in OSP presents new opportunities for enhancing the accuracy and efficiency of solutions, particularly in applications involving large datasets and incomplete measurements.

2.3.3 Case Studies in Placement Optimisation

Sensor placement optimisation has been widely applied in various fields, including structural health monitoring (SHM), power systems health monitoring, water distribution networks, non-destructive evaluation (NDE), condition-based maintenance (CBM), and prognostics and health management (PHM). In this section, an overview of some case studies that have implemented sensor placement optimisation techniques is given.

Structural health monitoring (SHM) is a field that aims to provide real-time information on the health condition of structures, ensuring their safety and preventing catastrophic failures. The use of sensor placement optimisation in SHM has been widely investigated in the literature. For example, Ostachowicz et al. [14] presented an unbiased state-of-the-art review of the research carried out in this area for researchers and practitioners in the SHM and optimisation fields. The review covers the definition of the optimisation problem, classification of techniques used, optimisation algorithms applied, and multi-objective optimisation. The authors of the reviewed article have focused on three commonly accepted and widely used methods in the SHM community, which are vibration-based monitoring, strain monitoring, and elastic wave-based monitoring.

Power systems are critical infrastructures that require constant monitoring to ensure their reliability and prevent power outages, also known as blackouts.

Sensor placement optimisation has been applied in power systems health monitoring to improve the accuracy and efficiency of fault detection and diagnosis. A Bayesian belief network (BBN) based approach has been proposed to optimise sensor placement for power systems health monitoring in the work by Pourali et al. [15]. The approach utilises functional topology, physical models of sensor information, and Bayesian inference techniques, along with constraints, to determine optimised sensor placement based on information metric functions. The methodology aims to address important questions such as inferring the health of a system or subsystem with limited monitoring points, using upward, downward, or distributed propagation techniques. The dynamic BBN serves as the engine for projecting the system's health. Such approaches are critical for ensuring the effective health monitoring of power systems while minimising costs associated with excessive sensor placement.

Water distribution networks are critical infrastructures that require constant monitoring to ensure their safety and prevent leaks and contamination. Sensor placement optimisation has been applied in water distribution networks to improve the efficiency and accuracy of leak detection and localisation. For example, Aral et al. [16] proposed a simulation-optimisation approach based on a single-objective function. The proposed model incorporates multiple factors used in the system's design to mimic a multi-objective approach, providing the final design without specifying a preference among the multiple objectives. A reliability constraint concept has also been introduced in the optimisation model to identify the minimum number of sensors and their optimal placement required to meet a pre-specified reliability criterion for the network. A progressive genetic algorithm approach has been utilised for the solution of the model by evolving subdomain sets of the complete set of junctions present in the system. The algorithm has been tested in two networks and compared with the outcome of other solutions presented in a water distribution systems analysis symposium, showing promising results for effective water sensor placement optimisation.

In a recent study, Kim et al. [17] discuss utilising the Convolution Neural Network (CNN) algorithm for NDE of aluminium panels. The objective is to classify the locations of defects by exciting the panel to generate ultrasonic Lamb waves, capturing the data through a sensor array, and then utilising deep learning to identify the features of 2D reflected waves from the defects. The study also explores the impact of optimal excitation location and sensor placement to improve the performance of the method. To ensure the training model's robustness and effective feature extraction, experimental data is collected by slightly varying the excitation frequency and defect location. The algorithm delivers high accuracy in classifying each defect location, even when a bar is attached to the panel.

PHM plays a crucial role in ensuring the safety and reliability of aerospace systems. Design for testability (DFT) is an important consideration for improving PHM performance, as information sensing and testing are the foundation of PHM. However, traditional DFT approaches, which only focus on fault detection and isolation requirements, are inadequate for sensor design and optimisation for PHM. To address this issue, a process for sensor selection and optimisation for PHM is proposed by Yang et al. [18]. Qualitative analysis of the intrinsic requirements of PHM for testability and quantitative definition of corresponding testability indexes are presented. Fault detection uncertainty is systematically analysed from various perspectives, including fault attributes, sensor attributes, and fault-sensor matching attributes. Object and constraint models for the sensor optimisation selection problem are studied in detail, and a sensor optimisation selection model is developed for aerospace system health management. The model considers the sensor total cost as the objective function and the proposed testability indexes under uncertainty test as constraint conditions. As the model is NP-hard, a genetic algorithm (GA) is introduced to obtain the optimal solution.

2.3.4 Cost function for Placement Optimisation

This section focuses on the significance of the cost function in guiding decision-making during sensor placement.

The cost function for sensor placement serves as a valuable tool in achieving an optimal sensor deployment that maximises system performance while considering the associated costs. By incorporating cost considerations into the placement optimisation process, engineers and decision-makers can make well-informed choices that align with budgetary constraints, ensuring cost-effectiveness without compromising system functionality.

A comprehensive cost function for sensor placement includes the following factors:

Sensor Coverage: Quantifying the extent to which sensors capture relevant information within the system's operational area. This includes assessing the spatial coverage and the quality of information gathered by the sensors. Sensor coverage can be quantified by assessing the spatial or temporal area that the sensors cover. This can be measured using metrics such as percentage coverage, spatial resolution, or time interval between data collection.

Sensor Connectivity: Evaluating the strength and reliability of sensor connections within the system. This encompasses metrics such as signal strength, connection success rate, and communication robustness. Sensor connectivity can be quantified by evaluating the ability of the sensors to establish and maintain reliable communication within the system. This can be measured based on metrics such as connection success rate, latency, or signal strength.

Interference Minimisation: Assessing the ability of the sensor placement to mitigate interference sources and maintain signal integrity. This includes considering interference rejection ratio, signal-to-interference ratio, and the effectiveness of interference mitigation techniques. Interference minimisation can be quantified by assessing the ability of the sensors to reduce or mitigate the impact of interference sources. This can be measured based on metrics such as signal-to-interference ratio, interference rejection ratio, or the ability to operate in noisy environments.

Resource Utilisation: Accounting for the efficient usage of system resources by the sensors, such as power consumption, bandwidth utilisation, and computational requirements. Resource utilisation can be quantified by evaluating how efficiently the sensors utilise system resources such as power, bandwidth, or processing capacity. This can be measured based on resource consumption rates or resource allocation efficiency.

Scalability can be quantified by assessing the ability of the sensor placement to accommodate system expansion or changes in the system's scale. This can be measured based on the ease of adding or removing sensors, as well as the impact on overall system performance.

These approaches provide a starting point for quantifying the objective functions in the Placement part of the sensor optimisation process. The first 4 factor is considered for the general cost function in the placement part; however, depending on the specific context and requirements of the complex engineering system, the actual quantification methods and metrics may vary.

$$\text{Cost}(f) = \alpha * \text{Sensor Coverage} + \beta * \text{Sensor Connectivity} + \gamma * \text{Interference Minimisation} + \delta * \text{Resource Utilization} \quad \textbf{(2-1)}$$

Quantifying the cost function for sensor placement involves assigning appropriate weights to each cost component based on their relative importance within the specific complex engineering system. These weights are determined through a thorough analysis, considering factors such as project budget, resource limitations, and the system's operational requirements. By employing multi-objective optimisation techniques, the cost function can be effectively integrated with other objective functions, such as sensor coverage, connectivity, and interference minimisation, to obtain an optimal sensor placement configuration.

By incorporating a well-defined cost function into the sensor placement optimisation process, complex engineering systems can achieve an optimal sensor deployment that maximises system performance while adhering to budget constraints. The cost-optimised sensor placement contributes significantly to the overall efficiency and success of the system.

2.4 Selection

Effective sensor selection can significantly improve the accuracy and reliability of health monitoring, fault diagnosis, and prognosis. For this reason, sensor selection is a challenging task that involves various considerations: the type and number of sensors, the performance characteristics of sensors, and the cost of sensor installation and maintenance.

This section presents an overview of sensor selection methods and techniques for IVHM applications. Starting with a discussion of the selection process and frameworks available in the literature, various methods for sensor selection, such as analytical, heuristic, and machine learning-based approaches, are explored. Selecting sensors based on information gain, such as mutual information-based, entropy-based, and Fisher information-based selection methods, as well as dynamic sensor selection optimisation perspectives, are also documented.

Lastly, evaluation of performance characteristics for sensor selection, including figures of merit, objective and cost functions, and information gain, is presented. The limitations and challenges of information gain optimisation, including complexity and computational requirements, a lack of comprehensive models and data, and the need for integration with other optimisation approaches and decision-making frameworks, are emphasised. Additionally, multi-objective optimisation and multi-criteria decision-making techniques for sensor selection optimisation, as well as sensor redundancy and implementation techniques, are presented.

2.4.1 Sensor Selection Methods

For clarity, it is essential to note that "sensor selection" has two distinct meanings in this field. The first refers to the process of selecting sensors from an existing network to optimise their performance by determining which sensors should be active at any given time. The second meaning, which is the focus of this work, pertains to selecting sensors for integration into a system during the design and build process. In this case, the selection process is geared towards choosing the most suitable sensors for the specific task at hand [2].

Several methods have been proposed to address sensor selection for IVHM systems, ranging from heuristic and analytical approaches to machine learning-based techniques. These methods differ in terms of their underlying assumptions, computational complexity, and solution quality. The selection process involves evaluating and comparing various sensor performance characteristics, including sensitivity, selectivity, reliability, and cost. Moreover, sensor selection should account for adaptation to changing operating conditions or system states. The first sensor selection use case was in aerospace systems, and it was based on design and performance requirements rather than a health management perspective. A model-based procedure that systematically selects an optimal sensor suite for overall health assessment of a designated host system is described in [19]. This procedure, known as the Systematic Sensor Selection Strategy (S4), was developed and implemented at NASA's John H. Glenn Research Centre with the primary objective of enhancing design phase planning and preparations for in-space propulsion health management systems.

Figure 2-5 illustrates the overall architecture of the S4 strategy, depicting the systematic approach to sensor selection and optimisation. Specifically, the S4 architecture operates through a two-phase filtering mechanism. The first phase is an iterative down-select process that evaluates a broad set of candidate sensor suites using an application-specific system diagnostic model and a sensor suite merit algorithm. This loop systematically discards inadequate combinations by weighing their fault-detection capabilities against predefined constraints such as cost, weight, or power. In the second phase, the shortlisted 'effective' sensor suites undergo a rigorous statistical evaluation algorithm to definitively identify the single optimal sensor suite for deployment. On the other hand, Figure 2-6 provides a detailed representation of the step-by-step process for applying the S4 strategy to a specific system.

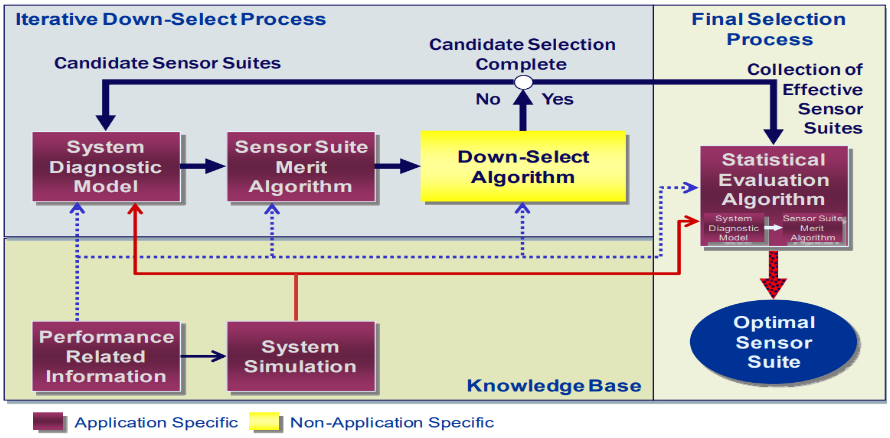

Figure 2-5 S4 Architecture [20]

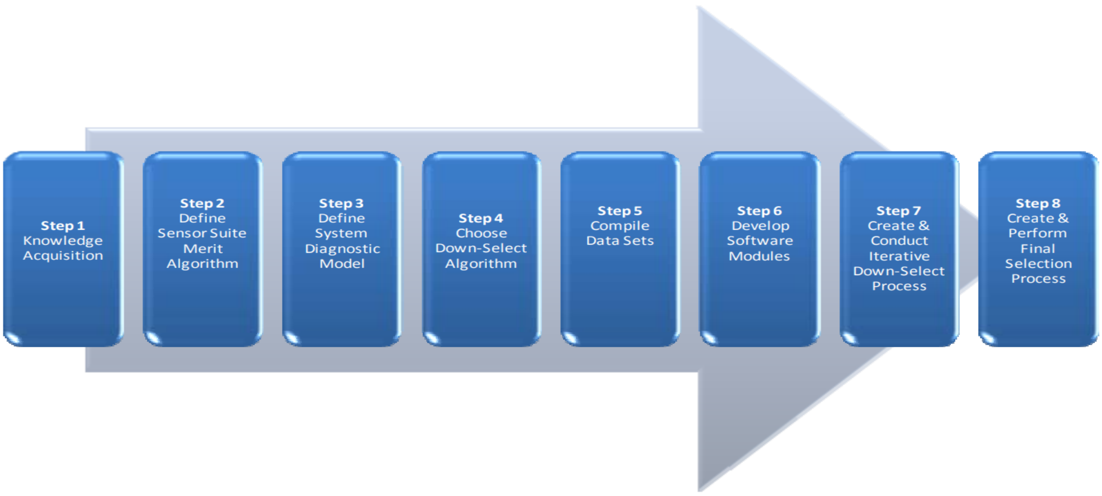

Figure 2-6 Process of applying the S4 strategy to a specific system [20]

The S4 strategy encompasses a well-structured and comprehensive framework for identifying and choosing the most suitable sensors to be integrated into in-space propulsion systems. Employing a systematic methodology ensures that the selected sensors are aligned with the system's requirements, performance objectives, and environmental conditions.

A general approach to sensor selection from a health management perspective was addressed in [20] with a proposed architecture that provides a justifiable, defendable sensor suite to address system health assessment requirements and additionally considers use outside of the aerospace community.

Sensor selection involves a multi-stage process that typically includes a knowledge-based selection, a down-select iteration, and a statistical evaluation algorithm. This process aims to identify the optimal set of sensors that can satisfy the requirements of the IVHM system while minimising the overall cost and complexity. The result of the analysis indicates that general sensor selection problems addressing diagnosability, or observability, are NP-complete and are therefore computationally intractable [21]. To solve complex problems in a reasonable time, approximate search solutions are needed. Brute Force or Exhaustive Search methods are ideal but can take too much time. Instead, refined search methods are used to find optimal or near-optimal solutions based on the objective function, which is an algorithmic representation of established figures of merits (FoMs) and system constraints.

Many techniques have been developed for general optimised solution searches. These optimisation techniques are well-documented in [22] Table 2-7 presents a selection of techniques identified in this literature review, along with their corresponding references.

Table 2-7 Sensor Selection Techniques Encountered in the literature review

Researcher	Technique	Application
Kulkarni et al. [2]	Proposes a method that utilises a scalable multi-objective framework for sensor selection to maximise fault detection rate while minimising the total cost of sensors	A wind turbine gearbox is considered to demonstrate the efficacy of the proposed framework.
Santi et al. [19]	A model-based procedure (S4) that systematically selects an optimal sensor suite for overall health assessment of a designated host system.	In-space propulsion health management systems
Maul et al. [20]	S4 was selected as the framework for further development and verification. Segmented into three groups: Knowledge Base, Iterative Down-Select Process, and Final Selection.	Applied to subsystem components of the Space Shuttle Main Engine
K. Nakai et al. [3]	Objective functions based on D-, A-, and E-optimality criteria of optimal design are adopted for greedy methods.	Applied to randomly generated systems and a practical dataset concerning the global climate.
Guan et al. [23]	Proposes a comprehensive evaluation method of sensor selection for PHM based on grey clustering.	Illustrated by an electronic control system, in which the effectiveness of different methods is compared
Joshi et al. [24]	Convex Optimisation	The problem of choosing sensors or measurements, from among a set of candidate measurements, to obtain the best resulting estimate of some parameters
Shamaiah et al. [25]	Greedy Sensor Selection Algorithm	The problem of sensor selection in resource-constrained sensor networks.
Wang et al. [26]	Entropy-based sensor selection method which can provide a quantitative description of the information contained in sensor data	Condition monitoring and prognostics of aircraft engines
Xu et al. [27]	Multi-objective Genetic Algorithm	Aircraft Engines
Najjar et al. [28]	The minimum Redundancy Maximum Relevance (mRMR) criterion with an unsupervised embedded algorithm	Heat Exchanger Fouling Diagnosis in Aerospace Systems
Jiao et al. [29]	Improved Binary Wolf Pack Algorithm	Typical discrete combinatorial optimisation problem
Manohar et al[30]	Balanced Model Reduction	Closed-loop Flow Control
Yan et al. [31]	Hybrid Bayesian Fisher information and mutual information	Unreliable sensor networks

The methodologies summarised in Table 2-7 established the foundational baseline for developing the MOSOF framework introduced in Chapter 3. Specifically, the inherent limitations identified within existing single-objective approaches underscored the critical need for the Pareto-based multi-objective optimisation strategy implemented in MOSOF. Furthermore, the entropy-based selection and S4 approach directly influenced the architectural decision to decouple the diagnostic index (NDCI) from the genetic algorithm. This separation ensures that the framework can efficiently evaluate severity-sensitive fault data without being computationally constrained by the optimisation loop. Additionally, the prevalence of greedy stepwise sensor selection in reviewed methods directly motivated MOSOF's own coverage-based subset construction, in which sensors are added iteratively in NDCI-rank order until a predefined diagnostic coverage threshold is met.

To solve an optimisation problem, it is essential to find a practical and achievable solution within a reasonable time frame. Heuristic approaches, such as genetic algorithms, particle swarm optimisation, and simulated annealing, rely on stochastic search strategies and can yield promising results. Machine learning-based methods, such as decision trees, random forests, and support vector machines, can also be effective in identifying the underlying patterns of a system by using a training dataset. One way to solve the sensor selection problem in the diagnostic domain is through information gain optimisation.

Several information gain-based methods have been proposed in the literature, such as mutual information-based methods, entropy-based methods, and Fisher information-based methods [26], [31]. These methods are effective in reducing the number of sensors required while maintaining high accuracy of health monitoring systems. While different sensor selection methods have their own advantages and limitations, it is crucial to consider the underlying assumptions, computational complexity, and solution quality of each method when selecting the appropriate method for a given application.

Finally, it is worth noting that prior knowledge and uncertainty should be taken into account in the sensor selection process. Previous knowledge about the system, its components, and their failure modes can provide valuable information for selecting the most informative sensors. Moreover, it is essential to consider the uncertainty associated with sensor measurements, the system model, and the diagnostic task when selecting sensors. Several studies have proposed methods for incorporating prior knowledge and uncertainty in the sensor selection process, such as Bayesian approaches [32] and fuzzy logic-based methods [33].

2.4.2 Evaluation of Performance Characteristics for Selection

Figure of Merits (FoMs) are quantitative measures of performance that can be used to evaluate the suitability of sensors for a particular application. FoMs provide a framework for assessing and comparing different sensors and selecting the best sensor for a given application. FoMs can be categorised into two groups: objective FoMs and subjective FoMs.

Objective FoMs are quantitative measures that can be directly calculated from the sensor data. They include sensitivity, specificity, accuracy, precision, and response time.

Subjective FoMs are qualitative measures that require expert judgment. Some examples of objective FoMs. They include ease of use, reliability, and maintainability.

Sensor selection methods in the diagnostic domain can be evaluated based on various performance characteristics, including fault detection rate (FDR), fault isolation rate (FIR), false alarm probability (FAP), and correct classification rate (CCR). These performance characteristics are influenced by the choice of sensors and their arrangement. Thus, it is essential to carefully evaluate the performance characteristics of sensors before selecting them for a specific application [34].

Objective functions are mathematical functions that need to be optimised to select the best sensors. These functions are designed to maximise the system's performance, subject to constraints such as budget, weight, and power consumption. Cost functions are also used to evaluate the cost of the sensor system and parameter trade-offs, including the cost of sensors, energy consumption, installation, maintenance, and replacement. The sensor selection problem requires achieving excellent performance while minimising costs. However, these two objectives often conflict as better performance typically comes with higher costs. Thus, researchers must find the optimal balance between cost and performance.

When setting up large, complex systems, it is essential to consider the cost of sensor configuration, as it often involves purchasing and installing a significant number of sensors. In addition to these upfront costs, ongoing usage costs should also be factored in, which may vary based on factors like connectivity, bandwidth, and sensor risk. Energy usage was the main objective for most of the sensor selection studies, in [34] The focus was solely on communication energy, which consists of detecting and transmission energy. There are other variables to consider, like in [35] The performance (fault detection reliability) of the sensor, as well as its cost (installation and communication), was also considered.

Information gain is a measure of how much a sensor measurement reduces the uncertainty about the system state. Information gain optimisation is a method of selecting sensors that maximises the amount of information gained from sensor measurements. It has been widely used in the design of sensor networks for various applications, including health monitoring, surveillance, and environmental monitoring. When implementing information gain optimisation in practice, several key considerations must be taken into account. These include cost constraints and trade-offs, sensor availability and compatibility, as well as robustness to noise and uncertainty. Fault trajectories, fault tolerance, fault detection, and fault isolation are essential considerations for the evaluation of information gain in diagnostics [19].

Although information gain optimisation is a powerful tool for sensor selection, it has several limitations and challenges. One of the main challenges is the complexity and computational requirements. Information gain optimisation involves calculating the mutual information between sensors and system states, which can be computationally expensive for large-scale systems. Another challenge is the lack of comprehensive models and data. Information gain optimisation requires accurate models of the system and the sensor characteristics, as well as precise data on the system states and sensor measurements. Finally, integration with other optimisation approaches and decision-making frameworks is also a challenge. Information gain optimisation needs to be integrated with other optimisation approaches and decision-making frameworks to ensure that the selected sensors meet the overall system requirements.

2.4.3 Multi-Objective Optimisation / Multi-criteria decision-making techniques for sensor selection optimisation

Sensor selection is often a multi-objective optimisation (MOO) problem that requires the simultaneous satisfaction of several objectives. These objectives can include diagnostic performance, robustness, and cost-effectiveness. In such cases, MOO techniques can be employed to consider the trade-offs between these objectives and to find the optimal sensor set that meets the desired criteria. In MOO, the goal is to find the Pareto optimal solutions. A solution is Pareto optimal if it represents the best possible trade-off between conflicting objectives. This set of Pareto optimal solutions forms the Pareto front, which means the boundary of the feasible solutions that cannot be further improved without sacrificing performance in other objectives. The evaluation of the Pareto optimal solutions can help decision-makers to identify trade-offs between different objectives and select the best solution based on their preferences and constraints.

Multi-criteria decision-making (MCDM) techniques, such as the Analytical Hierarchy Process (AHP) and the Technique for Order of Preference by Similarity to Ideal Solution (TOPSIS), can be used to rank sensors based on their diagnostic performance and other relevant criteria. AHP is a widely used MCDM technique that involves breaking down complex decisions into a hierarchy of objectives, criteria, and alternatives, and assigning relative weights to each level. TOPSIS is another MCDM technique that ranks alternatives by their similarity to the ideal solution and their distance to the worst solution.

Various MOO methods, including mathematical programming, evolutionary algorithms, and hybrid approaches, can be employed to solve the sensor selection problem. Mathematical programming methods, such as linear programming and mixed-integer programming, are efficient but may be limited by the complexity of the optimisation problem. Evolutionary algorithms, such as genetic algorithms [21], are robust and can handle complex optimisation problems, but may require a large number of iterations. Hybrid methods[36], which combine different optimisation techniques, can provide a balance between efficiency and robustness.

One popular method for multi-criteria decision-making is the Analytical Hierarchy Process (AHP), which involves decomposing the problem into a hierarchy of objectives, criteria, and alternatives. AHP then assigns weights to each element in the hierarchy and uses them to rank the other options based on their overall desirability [37]. Figure 2-7 illustrates a multi-level hierarchical structure for defining sensor criteria weights in fuel cell stack fault diagnosis.

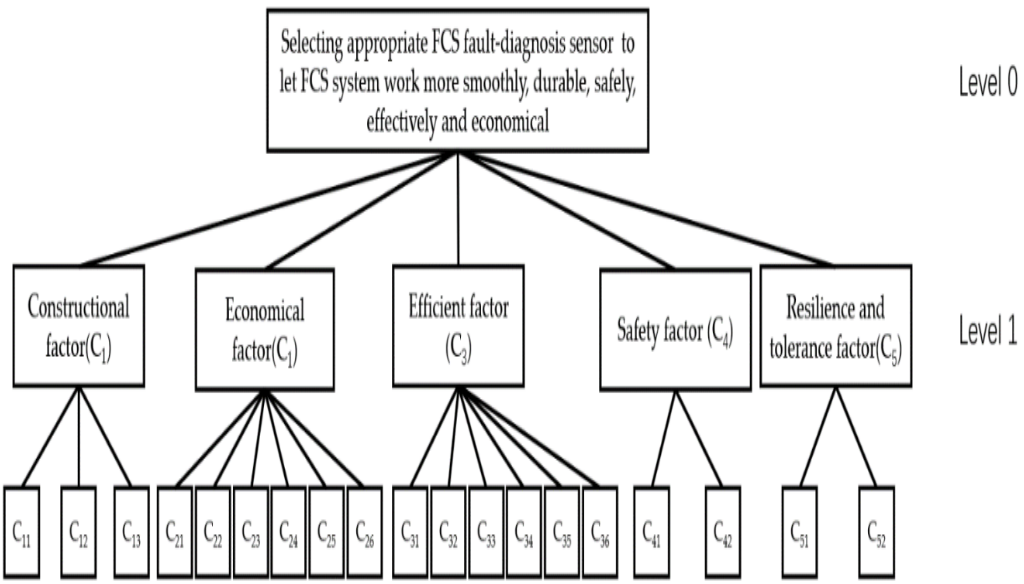

Figure 2-7 Multi-level hierarchical structure for fuel cell stack fault-diagnosis sensor criteria weight definition [24]

Another similar approach to AHP is grey clustering, a crucial technique for analysing and evaluating systems. Its efficiency in handling diverse feature types makes it indispensable in assessing sensor suite options. The recommended course of action is to assign a grey classification to the clustering object [23].

Another commonly used technique is TOPSIS, which involves ranking alternatives based on their similarity to an ideal solution and their distance to the worst solution. TOPSIS has been applied to sensor selection and MOO problems in several studies, [7], [38], [31].

The best method for sensor selection, whether it is MOO or MCDM, depends on the specific requirements and constraints of the application at hand. Each approach has its own set of advantages and limitations, and the choice should be made based on the characteristics of the problem and the goals of the sensor selection process.

The following factors can influence the decision on which method to choose:

Number of Objectives: If the sensor selection problem involves multiple, conflicting objectives, MOO is well-suited to identify trade-offs and find Pareto-optimal solutions. On the other hand, if the decision primarily revolves around a single dominant criterion, MCDM may be more appropriate.

Complexity of the Problem: For large-scale and complex sensor selection problems, MOO's ability to handle multiple objectives can be advantageous. However, if the problem is relatively simple and straightforward, MCDM may provide a quicker and more concise solution.

Decision-maker's Expertise: If the decision-maker is familiar with MOO techniques and capable of interpreting the Pareto front, MOO could be a compelling choice. Conversely, if simplicity and transparency are essential, MCDM may be the preferred approach.

Data Availability: The availability of reliable data for multiple objectives may influence the choice of method. If the data for different criteria are scarce or uncertain, MCDM might be a more practical option.

Ultimately, the best method for sensor selection will vary based on the specific context and the priorities of the decision-makers. In some cases, a combination of both MOO and MCDM might be used to leverage the strengths of each approach and arrive at a well-informed and balanced decision. Careful consideration of the problem's complexity, objectives, and available resources is crucial in making an informed choice for the sensor selection process.

2.4.4 Sensor Redundancy

Sensor redundancy is a crucial technique employed in various industrial and engineering applications to enhance the diagnostic reliability and fault tolerance of systems. By using redundant sensors to detect and isolate faults, the system's diagnostic reliability and accuracy can be improved. Once the system is built, its reliability remains constant and cannot be changed. However, what can be managed is the system's operational reliability or the ability to maintain its desired level of reliability over time.

Within the generic engineering design process (e.g., guided by standards such as ARP4754B), redundancy requirements are first identified during the Conceptual Design phase through a Functional Hazard Analysis (FHA). They are then formally allocated in the Preliminary Design phase via the Preliminary System Safety Assessment (PSSA) and finally implemented and verified during detailed design. Once hazardous or catastrophic failure conditions are identified, redundancy requirements (such as dual or triple-modular redundant sensor arrays) are allocated to specific subsystems to satisfy the required fault tolerance and probability boundaries before detailed hardware design commences.

This is where sensor redundancy comes into play. While the system’s inherent reliability cannot be changed, redundant sensors can be used to monitor the system's behaviour continuously and detect potential faults or deviations from the expected performance. By using redundant sensors, a system of checks and balances that helps ensure the system continues to operate within acceptable performance limits can be created.

The redundant sensors serve as a means of managing the system's operational reliability by providing real-time data on system health and performance. If a sensor malfunctions or provides erroneous readings due to a fault, the redundant sensors can act as backups and provide accurate data. This technique involves measuring the same parameter with multiple sensors, providing a backup in case of sensor failure or malfunction. Determining the optimal number of redundant sensors is crucial for achieving the desired level of performance while minimising the cost of the system. The optimal number of redundant sensors depends on several factors, such as system criticality, the likelihood of sensor failures, and cost considerations. Reliability-based approaches, probabilistic analysis, and cost-benefit analysis are some of the techniques used to determine the optimal number of redundant sensors. Developing algorithms capable of processing data from multiple sensors and selecting the correct measurement is necessary.

Implementing redundancy can be achieved through various approaches, with the voting algorithm being one of the most effective. This algorithm compares data from multiple sensors and selects the most accurate measurement, enhancing the system's diagnostic reliability and accuracy. Recent studies have proposed innovative algorithms and methodologies to improve the performance and availability of complex systems through sensor technology. One study [39] introduced a new voting algorithm that assigns priority to each sensor's measurement in real-time, allowing for accurate fault detection and isolation. Results from fault inoculation experiments demonstrated that the proposed algorithm outperforms the majority voter and the enhanced weighted average voter in terms of reliability and availability. While this algorithm is more complex and requires more comparisons and CPU time for each voting action, it can handle severe outliers and overcome the problem of having no clear majority, which exists in the majority voter. The traditional majority voting approach follows the majority decision, but it may struggle when there is no clear majority. The proposed voting algorithm, on the other hand, employs an adaptive prioritisation method to handle situations with no majority and make accurate decisions by giving higher importance to more reliable sensors in real-time. This approach enhances the reliability and availability of complex systems, making the voting algorithm more robust and effective in critical applications.

Another study [40] proposed a simple and efficient way to estimate the diagnostic coverage and false alarm values of redundant sensor systems using statistical methods. By evaluating these values, this methodology enables the development of a safety concept and functional safety analysis for sensor systems in safety-critical applications. This approach can also be helpful for optimising statistical sensor systems and ensuring the reliability of IVHM systems. In addition, [41] proposed statistical design, estimation, and optimisation approaches for efficient product definition and design of integrated sensor systems for safety-critical applications. This study highlights the limitations of relying on redundant configurations alone and proposes a methodology that optimises individual sensing channel performance and dependability figures, dependent on the redundant sensor output function and its diagnostic mechanism parameters. The proposed method was demonstrated for a redundant integrated linear Hall magnetic field sensor system in safety-critical automotive applications, providing a practical approach for sensor system architects to perform overall optimisation of redundant sensor systems, including dependability requirements.

Overall, to ensure that IVHM systems are reliable and readily available, it is essential to prioritise sensor redundancy. By utilising multiple sensors, the system's fault tolerance can be improved, and its diagnostic capabilities can be enhanced. After conducting a thorough review of the relevant literature, it is apparent that there are various strategies for optimising sensor redundancy in IVHM applications. Overall, these papers provide valuable insights into the optimisation of sensor redundancy for IVHM applications.

2.4.5 Cost Function for Selection Optimisation

This section explores the significance of the cost function in informing decision-making during sensor selection.

The cost function for sensor selection plays a pivotal role in achieving a balanced approach that maximises system performance while considering the associated costs. It allows engineers and decision-makers to weigh the trade-offs between sensor capabilities and their economic implications. By incorporating cost considerations into the optimisation process, a more cost-effective sensor configuration can be obtained, aligning with the project's budgetary constraints.

A comprehensive cost function for sensor selection includes the following factors:

Sensor Performance: Represents the performance characteristics of the sensors, including sensitivity, selectivity, reliability, accuracy, precision, response time, and robustness. In this context, selectivity is defined as the sensor's capacity to isolate a specific fault signature from background noise and confounding operational variables, effectively minimizing cross-sensitivity to unrelated parameters. Sensor performance can be quantified based on metrics such as accuracy, precision, sensitivity, response time, or any other relevant performance criteria.

System Compatibility: Refers to the compatibility of the selected sensors with the target system, considering factors such as physical dimensions, communication protocols, power requirements, environmental suitability, and integration complexity. Specifically, it denotes the sensor's physical and electronic integration capability within the existing aerospace environment. Practical examples of compatibility include adhering to strict electromagnetic interference shielding requirements, utilizing standardized avionics data bus protocols (e.g., ARINC 429), and meeting the rigorous physical footprint and thermal constraints of the designated mounting location. System compatibility can be quantified by assessing how well the sensors can integrate and communicate with other components of the complex engineering system. This can be measured based on compatibility protocols, communication standards, or successful integration tests.

Cost Effectiveness: Addresses the cost-related considerations in sensor selection. Cost effectiveness can be quantified by considering the total cost of ownership of the sensors, including acquisition costs, installation costs, maintenance costs, and any other associated expenses. A monetary value or a cost-to-benefit ratio can represent this.

Information Gain: Quantifies the amount of useful information that can be extracted from the sensor data for the intended application, considering factors such as data quality, relevance, comprehensiveness, and the potential for decision-making and analysis. Information gain can be quantified by evaluating how valuable the information the sensors can provide is for the system. This can be measured based on metrics such as data entropy reduction, information theory, or the ability to detect and identify relevant events or patterns.

Sensor durability can be quantified by assessing the expected lifespan, reliability, and robustness of the sensors under normal operating conditions. This can be measured in terms of mean time between failures (MTBF) or failure rates.

Sensor redundancy can be quantified by assessing the level of redundancy or backup sensors in the system. This can be measured based on the number of redundant sensors available or the ability to seamlessly switch between sensors in case of failure.

Sensor calibration stability can be quantified by assessing the ability of the sensors to maintain consistent and accurate calibration over time. This can be measured by evaluating calibration drift or the need for frequent recalibration.

Sensor interoperability can be quantified by evaluating the ability of the sensors to work seamlessly with other sensors and systems within the complex engineering system. This can be measured based on interoperability protocols, data exchange capabilities, or successful integration with other components.

These approaches provide a starting point for quantifying the objective functions in the Selection part of the sensor optimisation process. The first 4 factors are considered for the general cost function in the selection part; however, depending on the specific context and requirements of the complex engineering system, the actual quantification methods and metrics may vary.

$$\text{Cost (f)} = \alpha * \text{Sensor Performance} + \beta * \text{System Compatibility} + \gamma * \text{Cost Effectiveness} + \delta * \text{Information Gain} \tag{2-2}$$

To quantitatively represent the cost function (cost(f)), appropriate weights are assigned to each cost component based on their relative importance within the complex engineering system. These weights can be determined through rigorous cost-benefit analyses, considering factors such as the project budget, expected sensor lifespan, and specific operational requirements. By employing multi-objective optimisation techniques, the cost function can be effectively integrated with other objective functions, such as sensor performance and compatibility, to obtain an optimal and cost-effective sensor configuration.

By incorporating a well-defined cost function into the sensor selection process, complex engineering systems can make informed decisions that strike the right balance between sensor capabilities and economic considerations. The cost-optimised sensor configuration contributes significantly to the overall success and sustainability of the system.

2.5 Operation

The operation of sensor optimisation is a multifaceted process with specific objectives aimed at enhancing the overall performance and efficiency of sensor systems within complex applications. The operation of sensor optimisation is a multifaceted process with specific objectives aimed at enhancing the overall performance and efficiency of sensor systems within complex applications. The primary purpose of the operational side of sensor optimisation is to strategically select, configure, and calibrate the sensor in service to maximise its utility. This process involves several key objectives, such as data quality improvement, resource optimisation, robustness and fault tolerance, adaptability to changing conditions, and data fusion and integration.

The primary goal of the operation is to ensure that sensors are accurate and functioning properly. This can be subcategorised as: monitoring and control, maintenance optimisation, fault diagnosis and prognosis, and performance optimisation. It is imperative to capture the necessary information and adjust system parameters in response to changing conditions, thereby informing decision support systems through continuous monitoring. In summary, the operation of sensor optimisation serves the overarching purpose of improving the effectiveness and efficiency of sensor systems within complex applications.

2.5.1 Monitoring and Control

Real-time monitoring and control are critical aspects of complex systems. The use of sensors and data analytics enables real-time monitoring and control of systems, providing continuous feedback on the system's health and performance. By monitoring system parameters in real-time, faults can be detected and addressed before they escalate into costly system failures. In addition, real-time monitoring and control can also provide insight into system performance, enabling the optimisation of system parameters for increased efficiency and effectiveness.

Real-time monitoring and control are applicable to a wide range of systems, including manufacturing, transportation, and energy systems. In manufacturing, real-time monitoring and control can be used to optimise production processes, reducing waste and increasing throughput. In transportation, real-time monitoring and control can be used to optimise vehicle performance, reducing fuel consumption and improving safety. In energy systems, real-time monitoring and control can be used to optimise power generation and distribution, reducing costs and increasing reliability.

Zhou et al. [42] discussed the importance of condition monitoring (CM) in improving the reliability of rotating machinery (RM). They emphasised the need for an efficient CM method with simple and intuitive attributes for industrial applications. The development of health indicators (HIs) that connect fault detection, degradation assessment, and prognosis applications is crucial in CM. The paper reviews the construction methods of HIs for rotating machinery, covering both classical technical approaches and recent data-oriented intelligent methods such as deep learning. The benefits and potential of efficient HIs for condition monitoring are analysed, along with current challenges and future research opportunities.

By monitoring system parameters in real-time, safety-critical faults can be detected and addressed before they pose a threat to system operators or the public. In addition, real-time monitoring and control can be used to implement safety protocols, such as emergency shutdown procedures, in the event of a system failure.

Overall, real-time monitoring and control is a critical aspect of complex systems. By leveraging sensor data and advanced data analytics techniques, real-time monitoring and control can provide insight into system health and performance, enabling the optimisation of system parameters for increased efficiency and effectiveness. Real-time monitoring and control also have implications for system safety, enabling the detection and prevention of safety-critical faults.

2.5.2 Maintenance Optimisation

Maintenance optimisation is an essential aspect of complex systems, and it involves the use of various techniques to determine the optimal maintenance schedule for a system. The main goal of maintenance optimisation is to minimise the cost of maintenance while ensuring that the system operates at its optimal level. Two of the most commonly used maintenance optimisation techniques in IVHM systems are condition-based maintenance and predictive maintenance.

There are several factors to consider when selecting the appropriate maintenance optimisation technique for an IVHM system. These factors include the cost of maintenance, the system's criticality, the availability of replacement parts, and the system's operating environment. Additionally, it is important to select a maintenance optimisation technique that is compatible with the sensors and other monitoring tools used in the system.

In safety-critical systems, such as industrial plants or aircraft, the prevention of failure while maintaining high availability is crucial. Advanced prognostics algorithms and sensing techniques are being developed for predictive maintenance to achieve reliable and accurate predictions. However, there is a lack of in-depth studies on evaluating sensing techniques based on their prediction performance and inspection scheduling. Park et al. [43] addressed the need to evaluate the cost-effectiveness of different sensors by considering their contribution to reducing unnecessary inspection or measurement costs while maintaining prognosis performance. The authors conducted simulations to analyse prediction performance under varying measurement intervals and different levels of noise during degradation. Additionally, they analysed real run-to-fail (RTF) datasets from two different sensors to design an optimal measurement system for predictive maintenance. The study provides insights into selecting sensors based on cost-effectiveness and resistance to noise in order to improve maintenance strategies.

Demetriou et al. [44] introduced the economic aspect as a new factor in sensor selection for improved filtering of dynamical systems. The price of a single sensor, represented by high covariance values, is considered to incorporate the economic perspective into sensor optimisation for optimal filtering. Instead of relying on a single expensive and highly accurate sensor, the unit price and total price of a network of inexpensive, noisy sensors are utilised as alternatives. The study presents algorithms for integrated sensor optimisation for both finite and infinite-dimensional systems and provides examples to illustrate the effects of considering economic aspects in sensor selection.

Moradi et al. [45] addressed the challenging problem of performing and updating risk and reliability assessments for Complex Engineering Systems (CES) with high frequency. The complexity of operational data and system complexity necessitate the use of novel data-driven methods such as DL and engineering knowledge. The authors propose a mathematical architecture for operation condition and risk monitoring of CES, utilising a Bayesian Network (BN) to model system and subsystem relations, adverse event scenarios, and subsystem-level information fusion. Bayesian DL models are trained for subsystem diagnostics based on condition monitoring data, and their outputs are integrated into the BN. The proposed architecture effectively addresses both data and systems complexity, providing system-level insights and the ability to incorporate human inputs and qualitative information. The effectiveness of the approach is demonstrated through a case study on a Vapor Recovery Unit at an offshore oil production platform.

In conclusion, maintenance optimisation of complex systems, and the use of condition-based maintenance and predictive maintenance techniques can significantly improve the diagnostic reliability and performance of the system. The selection of the appropriate maintenance optimisation technique should be based on several factors, including the cost of maintenance, the criticality of the system, and the operating environment.

2.5.3 Fault Diagnosis and Prognosis

This section focuses on the techniques used for fault diagnosis and prognosis in complex systems. These techniques are essential for operation optimisation and system availability, as they allow for the detection and prediction of system faults before they can lead to system failure.

Gao et al. [46] focused on the growing need for early detection and identification of abnormalities and faults in industrial systems to minimise performance degradation and ensure safety. The authors highlight the importance of real-time fault diagnosis and fault-tolerant control methods in achieving these objectives. They provide a comprehensive review of fault diagnosis approaches and their applications, focusing on both model-based and signal-based perspectives. The paper aims to offer an extensive overview of the advancements in this field, with particular emphasis on the results reported in the last decade.

Baraldi et al. [38] proposed a general method for extracting a health indicator to measure the degradation state and predict the future evolution of industrial components. The method combines feature extraction techniques, including Empirical Mode Decomposition and Auto-Associative Kernel Regression, with a multi-objective Binary Differential Evolution (BDE) algorithm for optimal feature selection. The optimisation objectives focus on desired characteristics of the health indicator, such as monotonicity, trendability, and prognosability. A case study on turbofan engines is conducted to predict the remaining useful life. The results demonstrate the effectiveness of the proposed method in extracting accurate health indicators for prognostics.

2.5.4 Performance Optimisation

Performance optimisation is essential for efficient and effective IVHM, and sensor optimisation plays a necessary role in achieving this goal. IVHM systems need to be designed to optimise system performance, including operational efficiency, diagnostic reliability, availability, and maintainability. Performance optimisation aims to achieve these objectives by continuously monitoring and analysing system performance data and taking corrective actions when necessary.

Koutroulis et al. [47] reviewed the challenge of constructing comprehensive health indicators (HI) in prognostics and health management (PHM) using large amounts of condition monitoring data. The authors propose a novel anticausal-based framework with reduced model complexity to predict the cause from the effects of causal models, specifically designed for complex systems operating under time-varying conditions. Two heuristic methods, complexity estimation and Granger Causality, are used to infer the causal models. The framework demonstrates strong generalisation capabilities and robust online predictions of HIs, outperforming existing deep learning architectures in terms of average root-mean-square error (RMSE) by nearly 65%. The validation and comparison of the framework are conducted on NASA's N-CMAPSS dataset recorded from a commercial jet, further enhancing the CMAPSS simulation model.

Overall, performance optimisation techniques enable continuous monitoring and analysis of system performance data, allowing for real-time corrective actions to be taken to maintain optimal system performance. In conclusion, this section provides a comprehensive overview of sensor optimisation for improved operation in industrial systems. It covers various aspects, including real-time monitoring and control, maintenance optimisation, decision support systems, fault diagnosis and prognosis, and performance optimisation. The insights and strategies presented in this section contribute to the development of efficient and reliable sensor operation techniques in industrial applications.

2.5.5 Cost Function for Operation Optimisation

This section emphasises the significance of the cost function in guiding decision-making during sensor operation. The cost function for sensor operation plays a central role in achieving a reliable, energy-efficient, and optimised sensor performance while considering the associated costs. It allows engineers and decision-makers to strike a balance between system diagnostic reliability, energy consumption, maintenance requirements, and performance improvements while adhering to budget constraints.

A comprehensive cost function for sensor operation includes the following factors:

System Reliability: Evaluating the ability of sensors to perform their intended functions consistently and accurately over extended periods without failure or disruption, considering lifespan considerations. System reliability can be quantified by evaluating the probability of the system operating without failure over a given period. This can be measured using metrics such as mean time between failures (MTBF), mean time to repair (MTTR), or availability percentage.

Energy Efficiency: Assessing the energy consumption of sensors during operation to ensure optimal energy utilisation and reduce overall power consumption. Energy efficiency can be quantified by evaluating the energy consumption of the system in relation to the desired output or task. This can be measured based on metrics such as energy per unit of data processed, energy per unit of time, or energy efficiency ratings.

Maintenance Cost: Considering the costs associated with routine maintenance, sensor calibration, and periodic servicing to sustain the sensor's operational efficiency. Maintenance cost can be quantified by evaluating the expenses associated with maintaining and servicing the sensors and related components. This can be measured in terms of monetary costs, time required for maintenance activities, or the frequency of maintenance interventions.

Performance Optimisation: Quantifying the degree to which sensor operation aligns with the system's performance objectives, ensuring optimal utilisation of sensor data for decision-making. Performance optimisation can be quantified by evaluating the improvement in system performance achieved through optimisation efforts. This can be measured based on metrics specific to the system, such as throughput, accuracy, response time, error rates, or any other performance-related indicators.

These approaches provide a starting point for quantifying the objective functions in the operation part of the sensor optimisation process. Depending on the specific context and requirements of the complex engineering system, the actual quantification methods and metrics may vary.

Security can be quantified by evaluating the level of protection against unauthorised access, data breaches, or cyber threats. This can be measured based on metrics such as security vulnerability assessments, penetration testing results, or compliance with security standards.

Sensor Longevity can be quantified by evaluating the expected lifespan or operational duration of the sensors. This can be measured based on mean time between failures (MTBF), sensor degradation rates, or estimated lifetime usage.

The environmental impact of sensors can be quantified by evaluating the ecological footprint or sustainability aspects associated with their production, usage, and disposal. This can be measured based on metrics such as carbon footprint, material recyclability, or compliance with environmental regulations.

These approaches provide a starting point for quantifying the objective functions in the operation part of the sensor optimisation process. The first 4 factor is considered for the general cost function in the operation part; however, depending on the specific context and requirements of the complex engineering system, the actual quantification methods and metrics may vary.

$$\text{Cost}(f) = \alpha * \text{System Reliability} + \beta * \text{Energy Efficiency} + \gamma * \text{Maintenance Cost} + \delta * \text{Performance Optimisation} \quad \textbf{(2-3)}$$

Quantifying the cost function for sensor operation involves assigning appropriate weights to each cost component based on its relative importance within the specific complex engineering system. These weights are determined through comprehensive analysis, considering factors such as system requirements, maintenance schedules, energy budgets, and performance goals. By employing multi-objective optimisation techniques, the cost function can be effectively integrated with other objective functions, such as system reliability, energy efficiency, maintenance costs, and performance optimisation, to achieve an optimal sensor operation configuration.

By incorporating a well-defined cost function into the sensor operation optimisation process, complex engineering systems can achieve a reliable and energy-efficient sensor performance, optimising the overall efficiency and longevity of the system.

2.6 Data processing

Effective data processing is a critical component in the optimisation of sensor systems, aiming to maximise the information obtained. By employing advanced techniques, data processing enables the extraction of meaningful insights, leading to improved system performance and decision-making. This section explores various approaches and methodologies used in data processing to enhance the information gained from sensor data. The full potential of sensor data can be achieved by optimising the information gained from sensor data through effective signal processing, feature extraction and selection, machine learning techniques, and data fusion. In the following sections, the research delves into each of these approaches to understand their contributions in maximising the useful information obtained from sensor data.

2.6.1 Signal Processing Techniques

Signal processing techniques are employed to enhance the quality of sensor signals by reducing noise, amplifying relevant features, and mitigating interference. Through these techniques, engineers can maximise the information gained from sensor data, leading to improved system performance and more accurate analysis. The following techniques are discussed in this section: Filtering Techniques, Time-Frequency Analysis Techniques and Waveform Feature Extraction Techniques.

One of the most commonly used noise reduction techniques is filtering, which involves applying low-pass, high-pass, and band-pass filters to remove unwanted noise and artefacts from sensor signals. By selectively attenuating or amplifying specific frequency components, filtering improves the signal-to-noise ratio and enhances the quality of the acquired data. This allows for more precise analysis and interpretation of the sensor measurements. The Butterworth filter is a commonly used low-pass filter that removes high-frequency noise. Another widely utilised technique is the Kalman filter, which recursively estimates the system state based on sequential measurements, effectively filtering out noise. Additionally, adaptive filters are employed to adjust their parameters in response to changing noise conditions, allowing for the tracking of time-varying signals. An optimal filtering example of a stochastic singular system with correlated noises presented by Sun et al. [48], all results generalise the Kalman filtering, and its effectiveness is presented in a simulation example.

Time–frequency representation (TFR) has been a field of active research for the last few decades and remains a subject of interest today. A precise and accurate representation of nonstationary signals in the time–frequency domain is crucial, particularly in the context of mechanical fault diagnosis. Traditional TFRs depict the energy or power of signals in two-dimensional functions of time and frequency, effectively capturing fault signatures in diagnostic applications.

Various TFR methods employ different kernel functions, such as the short-time Fourier transform (STFT) with a linear kernel, the Wigner-Ville distribution (WVD) with a quadratic kernel, and the wavelet transform, which utilises an analysis basis constrained in both time and frequency.[9]

Waveform feature extraction techniques are employed to extract relevant features from sensor data. Peak detection is a commonly used technique that identifies the maximum or minimum values within a waveform, providing insights into signal characteristics [49]. Zero-crossing detection identifies the points at which a waveform crosses the horizontal axis, offering information about signal behaviour. Fourier Analysis, which decomposes a signal into its frequency components using the Fourier transform, enables further analysis and interpretation for larger datasets [50].

In summary, signal processing techniques are essential for optimising the information gained from sensor data. By effectively applying these signal processing techniques, engineers can extract valuable insights, improve system performance, and make informed decisions based on the processed sensor data.

2.6.2 Feature Extraction and Selection

Feature extraction and selection methods are utilised to identify and extract relevant information from the processed sensor data. By focusing on key features, these techniques reduce the dimensionality of the data and the computational burden associated with processing large datasets, thereby enhancing computational efficiency and highlighting the most informative aspects.

Feature extraction is an important step in the data processing phase of the basic condition monitoring process. This extraction process is particularly important for handling noisy sensor data and avoiding excessive input features, especially in the case of vibration data, during the classifier learning phase. Therefore, feature extraction is often considered the first and essential step in any classification task [51].

Standard basic features include maximum, mean, minimum, peak, peak-to-peak interval, and others. Additionally, more complex feature extraction methods like principal component analysis (PCA), independent component analysis (ICA), and kernel principal component analysis (KPCA) can be employed [52]. These advanced methods enable the extraction of more intricate and informative features from the sensor data, enhancing the accuracy and effectiveness of classification algorithms.

Feature selection has gained significant attention in recent years in machine learning applications. Its objective is to identify and retain the most relevant features from an original dataset, aiming to enhance the quality and efficiency of feature sets used in various tasks such as classification, regression, clustering, and time-series prediction. This objective can be achieved using a variety of methods, including filter, wrapper, and embedded approaches.

Irrelevant or redundant features can lead to overfitting and performance degradation, making feature selection essential for mitigating these issues. By reducing the dimensionality of the dataset and selecting the most informative subset of features, feature selection techniques offer benefits such as improved interpretability of models, reduced computational costs, and enhanced learning accuracy. These techniques have found widespread adoption across domains like text mining, image analysis, and biomedical research. The visual representation of the feature selection process, as depicted in Figure 2-8, showcases the transformation of an original feature set into a carefully selected subset, resulting in improved performance and efficiency of machine learning algorithms [53].

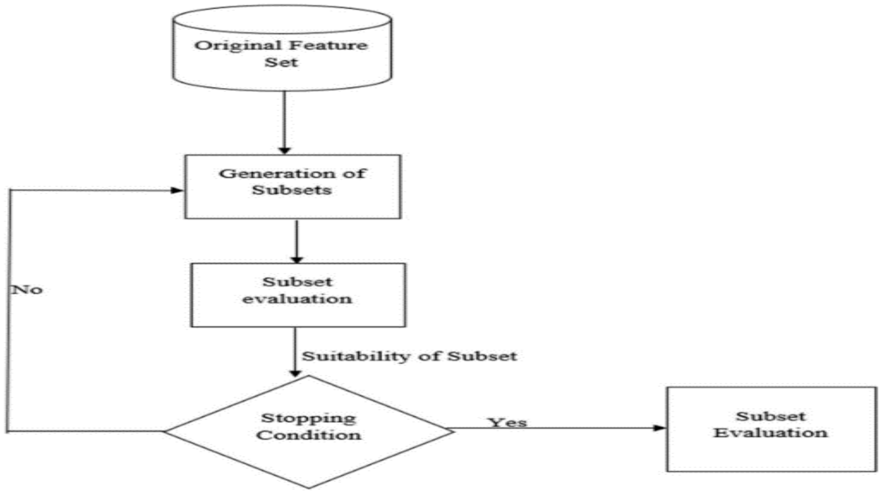

Figure 2-8 Illustration of the general feature selection process [50]

Filter approaches involve ranking features based on statistical measures such as correlation or mutual information. Wrapper approaches evaluate the performance of a specific machine learning algorithm using different subsets of features and select the subset that yields the best performance. Embedded approaches involve selecting features as part of the training process for the machine learning algorithm [54].

Several feature extraction and selection techniques have been applied in complex systems, each with its own strengths and limitations. In a study on feature selection for pattern classification systems, Peng et al. [55] investigated the use of the minimal-redundancy-maximal-relevance criterion (mRMR) based on mutual information. Their goal was to select a compact set of superior features at a low cost. The proposed approach involved a two-stage feature selection algorithm that combined mRMR with other advanced feature selectors, such as wrappers. The algorithm was extensively evaluated using different classifiers (naive Bayes, support vector machine, and linear discriminant analysis) and diverse datasets (handwritten digits, arrhythmia, NCI cancer cell lines, and lymphoma tissues). The experimental results demonstrated the promising improvement in feature selection and classification accuracy achieved by incorporating mRMR.

The choice of feature extraction and selection technique depends on various factors, including the specific problem being addressed, the type of data being analysed, and the available computational resources. It is crucial to evaluate the performance of different techniques in the context of the given application to determine the most suitable technique.

2.6.3 Machine Learning Techniques

Machine learning (ML) techniques have emerged as powerful tools for data processing in sensor systems. By leveraging algorithms such as classification, regression, and anomaly detection, machine learning can automatically learn patterns and relationships within sensor data. This enables the system to make predictions, detect anomalies, and uncover complex insights that might not be immediately apparent. Integrating machine learning into data processing enhances the system's ability to extract valuable information and optimise the sensor system's performance.

In the field of PHM, data-driven methods, particularly ML and deep learning (DL) techniques, have gained widespread adoption for tasks such as anomaly detection, fault diagnostics, and prognostics [45]. These methods possess the capability to handle large volumes of highly non-linear data effectively. DL models excel at processing operational data and automatically generating features for various tasks, including detection, classification, and prediction of patterns within the data. This reduces the reliance on domain expertise and extensive manual feature engineering, particularly when complete and representative data is available.

Learning problems in the context of DL can be categorised into four main groups: Supervised, Unsupervised, Semi-supervised, and Reinforcement learning. To implement DL algorithms, three key components are required: (1) training and testing data, (2) an objective function, and (3) an optimisation scheme. Variations in these components give rise to a multitude of distinct DL algorithms and architectures, such as Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs), Long Short-Term Memory Networks (LSTMs), Generative Adversarial Networks (GANs), and Autoencoders [45].

Rezaeianjouybari et al. [56] presented the state of the art, challenges, and opportunities for deep learning applications in PHM. Well-documented lists of activation functions in deep learning, optimisation algorithms in deep architectures. In the context of data-driven predictive maintenance (PdM), the use of machine learning (ML) algorithms has gained attention for their potential in artificial intelligence. However, the performance of ML algorithms can be compromised when dealing with high-dimensional and discontinuous machine data. Standard dimension reduction techniques may not effectively handle such challenges. To address this, Aremu et al. [57] proposed an ML-based dimension reduction framework that clusters observations based on data modality and utilises Laplacian eigenmaps embedding to obtain low-dimensional representations. The framework is applied to the Commercial Modular Aero-Propulsion System Simulation dataset, demonstrating its effectiveness in handling high-dimensional discontinuous machine data for ML-based PdM analysis.

Thoppil et al. [58] discussed popular deep learning architectures and their significance in machinery health prognostics, using benchmark time-series machinery failure datasets, and highlighting the contributions of researchers in implementing deep learning approaches for accurate machinery health diagnostics and prognostics. Saufi et al. [59] discussed deep learning models and their applications in machinery fault detection and presented classification tables for the models used in different diagnosis stages, such as fault detection, fault identification, fault size estimation and fault growth prediction.

Intelligent fault diagnosis (IFD) has gained attention for automating machine fault recognition and reducing human labour. However, existing reviews lack comprehensive coverage and provide limited guidance for future studies. To address this, Lei et al. [60] provided a systematic review and roadmap of IFD's development, encompassing traditional machine learning theories, the advent of deep learning, and the prospects of transfer learning. The roadmap highlights potential research trends and challenges in IFD. Yang et al. [61] also presented a detailed textbook about the foundation of transfer learning and its applications.

ML techniques serve as the diagnostic classification layer in MOSOF, consuming NDCI-optimised sensor data. Their selection, configuration, and unbiased evaluation via repeated nested cross-validation within this framework are detailed in Chapters 3 and 4. Overall, the choice of ML technique depends on the specific application and the characteristics of the data. It is important to carefully evaluate different techniques and select the most appropriate one for the problem at hand.

2.6.4 Data Fusion Techniques

Data fusion techniques are widely used in complex systems to improve the accuracy, reliability, and robustness of the results by integrating information from multiple sources. Data fusion can be performed at different levels, such as sensor-level fusion, feature-level fusion, and decision-level fusion, depending on the application requirements and the available data. Figure 2-9 represents three levels of information fusion techniques used in diagnostic and decision support systems. The selection of the appropriate fusion technique depends on the characteristics of the data, the desired outcome, and the computational resources available.

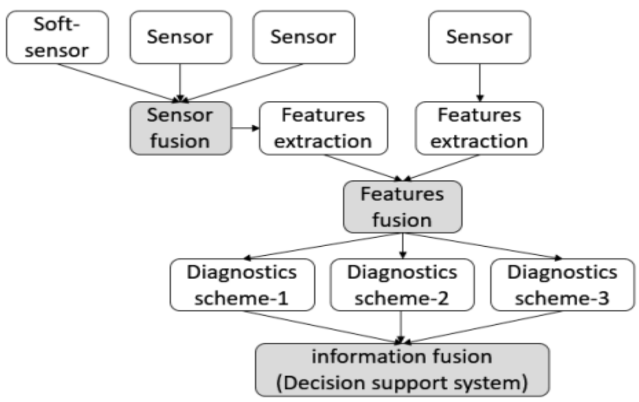

Figure 2-9 Three levels of information fusion for diagnostic and decision support systems

In the realm of sensor data management, the efficient handling of data from multiple sources is paramount. Sensor-level fusion combines the raw data from multiple sensors (multi-source sensors) to form a more accurate and reliable representation of the underlying physical phenomenon. It is particularly useful in situations where the sensors have different response characteristics or are affected by different sources of noise. Sensor-level fusion can be achieved by techniques that range from simple weighted averages to more advanced methods such as fuzzy logic, Kalman filters, and probabilistic approaches. Several studies have demonstrated the effectiveness of sensor data fusion in various applications [62].

Feature-level fusion combines the extracted features from multiple sensors to form a more informative and discriminative representation of the underlying physical phenomenon. Feature-level fusion is particularly useful in situations where the sensors have different measurement modalities or provide complementary information. In the context of vibration monitoring, feature-level fusion has been successfully applied for sensor fault detection, combining correlated measurements through statistical feature extraction [63].

At the decision level of the diagnostic and prognostic process, the fusion of results from multiple independent methods can enhance the accuracy and confidence of estimations. Different techniques may be more effective in identifying specific problems, making their combination valuable. Decision-level fusion can be achieved by voting, weighting, or selecting the most reliable algorithm or model based on its performance and the uncertainty of the results. It is particularly useful in situations where the algorithms or models have different assumptions or are trained on different data. Additionally, estimations from these methods can be combined with other sources of information, such as vibration analysis, maintenance history, observations during inspections, and negative information. These fusion approaches occur at two levels: automated decision level and supervised decision level. While significant research has been conducted on the automated decision level, there is a scarcity of examples in the literature regarding the supervised decision level [64].

Overall, in a self-organised distributed system, collective decision-making plays a crucial role. Various approaches have been proposed for collective decision-making, including voting models, swarm methods inspired by biology, and methods for task and role allocation. However, decentralised information fusion systems face a significant challenge in terms of real-time communication. Current real protocols rely on a central control unit to manage communication timing and flow. Overcoming this challenge is essential for the development of future approaches to decentralised information fusion systems [65].

2.6.5 Cost function for Data Processing Optimisation

This section highlights the significance of the cost function in guiding decision-making during data processing.

A comprehensive cost function for data processing includes the following factors:

Data Accuracy: Quantifying the level of agreement between sensor measurements and ground truth values, ensuring that the processed data is reliable for decision-making. Data accuracy can be quantified by assessing the level of agreement between sensor measurements and ground truth values or reference data. This can be measured using metrics such as mean absolute error, root mean square error, or statistical measures of accuracy.

Computational Efficiency: Evaluating the efficiency of data processing algorithms and techniques to minimise resource consumption, such as processing time, memory usage, and energy consumption. Computational efficiency can be quantified by evaluating the computational resources required for processing sensor data. This can be measured based on metrics such as processing time, memory usage, or energy consumption during data processing.

Feature Relevance: Assessing the significance of extracted features from sensor data to ensure that the most relevant and informative features are utilised for analysis. Feature relevance can be quantified by assessing the significance of extracted features from sensor data for the intended analysis or decision-making process. This can be measured based on metrics such as feature importance scores, information gain, or correlation coefficients.

Model Complexity: Considering the complexity of data processing models, it is necessary to strike a balance between accuracy and computational overhead, avoiding overly complex models that might be resource intensive. Model complexity can be quantified by evaluating the complexity or simplicity of the models used for data processing. This can be measured based on metrics such as the number of parameters, the depth of the model, or the computational complexity of the algorithms.

Sensor Data Integration can be quantified by assessing the ability to combine and merge data from multiple sensors to create a comprehensive view of the system. This can be measured based on the effectiveness of data fusion algorithms, data alignment accuracy, or the quality of integrated data outputs.

Sensor Data Privacy can be quantified by evaluating the level of protection and confidentiality applied to sensor data. This can be measured based on privacy-preserving techniques, encryption methods, or compliance with data privacy regulations.

These approaches provide a starting point for quantifying the objective functions in the Data Processing part of the sensor optimisation process. The first 4 factors are considered for the general cost function in the data processing part; however, depending on the specific context and requirements of the complex engineering system, the actual quantification methods and metrics may vary.

$$\text{Cost}(f) = \alpha * \text{Data Accuracy} + \beta * \text{Computational Efficiency} + \gamma * \text{Feature Relevance} + \delta * \text{Model Complexity} \tag{2-4}$$

Quantifying the cost function for data processing involves assigning appropriate weights to each cost component based on its relative importance within the specific complex engineering system. These weights are determined through comprehensive evaluations, considering factors such as data processing performance requirements, available computational resources, and budget constraints. By employing multi-objective optimisation techniques, the cost function can be effectively integrated with other objective functions, such as data accuracy, computational efficiency, feature relevance, and model complexity, to achieve an optimal data processing configuration.

By incorporating a well-defined cost function into the data processing optimisation process, complex engineering systems can achieve an efficient and cost-effective utilisation of sensor data. The cost-optimised data processing contributes significantly to the overall efficiency and decision-making capabilities of the system.

2.7 Conclusion

In this comprehensive literature review, the various aspects of sensor optimisation in complex engineering systems are explored. The review encompassed sensor selection, placement, data processing, and operation, each of which plays a crucial role in enhancing system performance, reliability, and cost-effectiveness.

Throughout the review, the importance of cost functions as indispensable tools for guiding decision-making in each stage of sensor optimisation is highlighted. The cost functions allowed engineers and decision-makers to strike a balance between performance requirements and financial considerations, ultimately leading to cost-effective and efficient sensor configurations.

In the sensor placement stage (Section 2.3), the cost function facilitated optimal sensor deployment by accounting for coverage, connectivity, interference minimisation, and resource utilisation costs. Through comprehensive analysis, engineers achieved sensor placements that maximised system performance while minimising operational costs.

Similarly, in the sensor selection stage (Section 2.4), how the cost function played a pivotal role in evaluating the economic implications of sensor choices is discussed. By considering acquisition costs, installation expenses, maintenance requirements, and other relevant factors, the cost function enabled the selection of sensors that aligned with project budgets while meeting system performance objectives.

In sensor operation (Section 2.5), the cost function guided decisions to achieve a reliable, energy-efficient, and cost-effective sensor performance. By considering system reliability, energy consumption, maintenance costs, and performance optimisation, the cost function optimised sensor operations to support the system's long-term sustainability and performance excellence.

In data processing (Section 2.6), the cost function played a vital role in ensuring an efficient analysis and utilisation of sensor data. By balancing data accuracy, computational efficiency, feature relevance, and model complexity, the cost function enabled the extraction of valuable insights while minimising computational overhead and resource utilisation.

To systematically quantify these differing priorities, a weighting matrix was developed using a 1-3-9 scale—a standard approach derived from quality function deployment and AHP methodologies. In this scale, 1 denotes low priority, 3 denotes moderate priority, and 9 denotes critical priority. Table 2-8 outlines these differing assumptions across 25 cost function aspects.

Table 2-8 Stakeholder Prioritisation Matrix for Sensor Optimisation Factors

Cost Functions Aspect	***Most Useful Quantification Approach***	***OEM***	***Airline***	***MRO***
Sensor Performance	Accuracy, Precision	3	9	9
System Compatibility	Integration Complexity, Communication Protocols	9	3	9
Cost-Effectiveness	Total Cost of Ownership, Cost-to-Benefit Ratio	9	9	1
Information Gain	Data Entropy Reduction, Information Theory	3	9	9
Sensor Durability	MTBF, Failure Rates	9	9	3
Sensor Redundancy	Number of Redundant Sensors, Seamless Switching Capability	9	9	3
Calibration Stability	Calibration Drift, Frequency of Recalibration	3	9	9

Sensor Interoperability	Data Exchange Capabilities, Successful Integration with Other Components	9	3	9
Sensor Coverage	Percentage Coverage, Spatial Resolution	9	9	9
Sensor Connectivity	Connection Success Rate, Latency	9	3	9
Interference Minimisation	Signal-to-Interference Ratio, Interference Rejection Ratio	9	3	3
Resource Utilisation	Power Consumption, Bandwidth Utilisation	9	9	1
Data Accuracy	Mean Absolute Error, Root-Mean-Square Error	9	9	9
Computational Efficiency	Processing Time, Memory Usage	9	3	3
Feature Relevance	Feature Importance Scores, Information Gain	3	9	9
Model Complexity	Number of Parameters, Model Depth	9	3	3
Sensor Data Integration	Data Fusion Effectiveness, Data Alignment Accuracy	9	9	9
Sensor Data Privacy	Encryption Methods, Data Privacy Regulations	9	9	3
System Reliability	MTBF, MTTR	9	9	9
Energy Efficiency	Energy Per Unit of Data Processed	9	3	3
Maintenance Cost	Monetary Costs, Time Required for Maintenance Activities	3	9	9
Performance Optimisation	Throughput, Accuracy	9	9	9
Security	Security Vulnerability Assessments, Penetration-Testing	9	9	3
Sensor Longevity	MTBF, Sensor Degradation Rates	9	9	9
Environmental Impact	Carbon Footprint, Material Recyclability	9	3	3

To illustrate how the diverse factors in Table 2-8 feed into the consolidated master equations, Figure 2-10 presents a mind-map of the cost function relationships of all the aspects and their logical flow.

Figure 2-10 Mind-map illustrating the relationship between optimisation aspects and stakeholder requirements

Figure 2-11 visually captures the resulting AHP weighting profile, demonstrating how the distinct operational goals of the OEM, Airline, and MRO are translated into explicit mathematical constraints. By converting the qualitative 1-3-9 scores from Table 2-8 through AHP pairwise comparisons, the profile reveals the dominant drivers for each stakeholder. For example, the MRO's weighting profile is heavily skewed toward diagnostic accuracy and feature relevance, whereas the Airline's profile demands a balance heavily weighted by sensor durability and calibration stability to ensure continuous fleet availability. Meanwhile, the OEM profile strongly favours system compatibility, initial cost-effectiveness, and physical resource utilisation. Ultimately, these mathematically normalised coefficients depicted in the profile are based on the predicted values injected into the master fitness function. These values must be evaluated based on real world applications requirements, ensuring that the sensor optimisation algorithm explicitly solves for the designated stakeholder's operational reality.

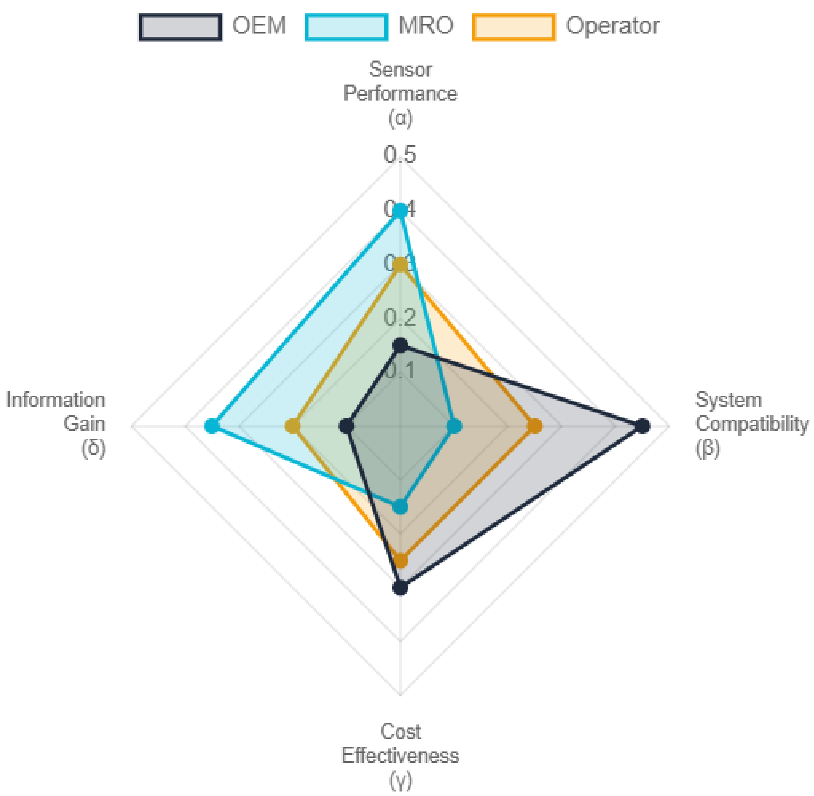

Figure 2-11 AHP Weighting Profile

In conclusion, this literature review demonstrated the significant role that cost functions play in the optimisation of sensors in complex engineering systems. By considering cost implications alongside other performance metrics, cost functions enable well-informed decisions that lead to the success and sustainability of the system. As sensor technology continues to advance, the development and refinement of cost functions will remain essential in driving innovation and efficiency in the field of sensor optimisation.

2.7.1 Summary of Key Findings

Throughout this literature review on sensor optimisation in complex engineering systems, several key findings have emerged, shedding light on the crucial role of cost functions and their impact on decision-making. The review examined sensor placement, selection, operation and data processing, as well as an integrated approach for multi-objective sensor optimisation. In this section, the main findings that contribute to a deeper understanding of the significance of cost functions in sensor optimisation are summarised:

Balancing Performance and Cost: Cost functions provide a systematic approach to balancing sensor performance requirements with associated costs. By quantifying and weighing cost components, engineers and decision-makers can make informed choices that optimise system performance while adhering to budgetary constraints.

Optimising Sensor Placement Efficiency: In the context of sensor placement, the cost function guides decisions to achieve optimal sensor coverage, connectivity, interference minimisation, and resource utilisation. By incorporating cost considerations, engineers achieve cost-efficient sensor deployments without compromising system performance.

Cost-Effectiveness in Sensor Selection: The cost function for sensor selection considers various cost-related factors, such as sensor acquisition costs, installation expenses, maintenance, and long-term operational costs. By evaluating these components, decision-makers can select sensors that meet performance criteria while remaining cost-effective.

Reliable and Energy-Efficient Sensor Operation: Cost functions play a vital role in optimising sensor operation, balancing system reliability, energy efficiency, maintenance costs, and performance improvements. By quantifying these cost components, decision-makers achieve reliable and energy-efficient sensor performance.

Efficient Data Processing Strategies: The cost function for data processing ensures the efficient utilisation of sensor data. By considering data accuracy, computational efficiency, feature relevance, and model complexity, engineers can extract valuable insights while minimising computational overhead and resource consumption.

Integrated Multi-Objective Optimisation: The integrated approach to multi-objective sensor optimisation employs cost functions alongside other objective functions, resulting in balanced and efficient sensor configurations. Trade-off analysis and Pareto front analysis enable decision-makers to make informed choices that optimise multiple objectives simultaneously.

Role of Cost Functions in Future Sensor Optimisation: As technology evolves and sensor optimisation advances, cost functions will continue to be critical tools in guiding decision-making. By considering cost implications alongside other performance metrics, cost functions support innovation and efficiency in sensor optimisation.

In conclusion, this literature review underscores the vital role of cost functions in sensor optimisation for complex engineering systems. By effectively incorporating cost considerations throughout the sensor optimisation process, engineers and decision-makers can achieve cost-effective, efficient, and reliable sensor configurations. As the field of sensor technology continues to evolve, cost functions will remain indispensable in driving advancements and ensuring the successful implementation of optimised sensor solutions in diverse engineering applications.

This literature review makes several key contributions to the understanding of sensor optimisation in complex engineering systems:

Comprehensive Overview: The review provides a comprehensive overview of sensor optimisation, covering essential aspects such as sensor selection, placement, data processing, and operation. By examining each stage of the optimisation process, the review offers a holistic perspective on the challenges and opportunities in achieving efficient and cost-effective sensor configurations. A taxonomy and a concept map of the area have been generated.

Role of Cost Functions: The review emphasises the critical role of cost functions in guiding decision-making at each stage of sensor optimisation. By quantifying the trade-offs between performance and cost, cost functions empower engineers and decision-makers to make informed choices that align with budgetary constraints while optimising system performance.

Integration of Objective Functions: The integrated approach to multi-objective sensor optimisation highlighted in the review illustrates the importance of considering all objective functions simultaneously. By combining cost functions with other performance metrics, decision-makers can achieve balanced and efficient sensor configurations that cater to multiple optimisation goals.

2.7.2 Implications

The findings of this literature review have several implications for research, engineering practice, and decision-making in the field of sensor optimisation:

Informed Decision-Making: By incorporating cost functions into the optimisation process, decision-makers can make informed choices that balance system performance requirements with financial considerations. Cost functions provide a quantitative basis for decision-making, promoting cost-effective and efficient sensor configurations.

Performance-Driven Optimisation: The integration of cost functions with other objective functions emphasises the need for performance-driven optimisation. Decision-makers must consider not only cost implications but also overall system performance to achieve successful sensor configurations.

Future Research Directions: The review identifies emerging trends and future research directions, including advancements in sensor technology and the integration of AI and machine learning. Researchers can use this information as a foundation to explore innovative approaches to sensor optimisation.

Based on the insights gained from this review, decision-makers involved in sensor optimisation in complex engineering systems are recommended to:

Integrate Cost Functions: Incorporate cost functions into the decision-making process at each stage of sensor optimisation. By quantifying cost implications, decision-makers can make cost-effective choices that align with system performance objectives.

Consider Multi-Objective Optimisation: Embrace an integrated approach to multi-objective optimisation, combining cost functions with other performance metrics. This ensures a balanced and efficient sensor configuration that meets multiple optimisation goals.

Monitor Emerging Trends: Stay informed about emerging trends and advancements in sensor technology. Continuously evaluate how these advancements can improve sensor optimisation practices and enhance system performance.

In conclusion, this literature review highlights the significant role of cost functions in guiding efficient and cost-effective sensor optimisation in complex engineering systems. By considering cost implications alongside other performance metrics, engineers and decision-makers can achieve optimised sensor configurations that support the success and sustainability of complex engineering applications. The review's findings offer valuable insights for researchers, engineers, and decision-makers seeking to enhance sensor optimisation practices and improve decision-making in diverse engineering scenarios.

2.7.3 Recommendations for Future Research

This literature review has provided valuable insights into the role of cost functions and their impact on decision-making. Building on these findings, several areas for future research are recommended to further advance the field of sensor optimisation:

Enhancing Cost Function Models: Future research can focus on refining and expanding cost function models for sensor optimisation. This includes investigating advanced techniques for quantifying cost components and exploring novel approaches to weigh the importance of different cost factors in specific engineering applications.

Dynamic Cost Functions: Investigate the development of dynamic cost functions that adapt to changing operational conditions and evolving project budgets. Dynamic cost functions can provide real-time decision support, enabling sensor configurations that respond to varying performance requirements and cost constraints.

Uncertainty and Risk Analysis: Incorporate uncertainty and risk analysis into cost functions to account for uncertainties associated with sensor performance, cost estimations, and environmental variations. Understanding the impact of uncertainties on optimisation outcomes can lead to more robust and reliable sensor configurations.

Integration of Lifecycle Cost Analysis: Extend cost functions to include lifecycle cost analysis, considering the long-term costs associated with sensor maintenance, calibration, and replacement. Lifecycle cost analysis can provide a comprehensive view of cost implications over the sensor's entire operational lifespan.

Optimisation Algorithms: Explore advanced optimisation algorithms that effectively handle the multi-objective nature of sensor optimisation. Investigate the application of evolutionary algorithms, genetic algorithms, and machine learning techniques to efficiently search for optimal sensor configurations within the multi-dimensional cost-performance space.

Real-World Deployment Studies: Conduct real-world deployment studies to validate the effectiveness of cost function-driven sensor optimisation in practical engineering applications. These studies can provide valuable insights into the challenges and benefits of implementing cost-optimised sensor configurations in diverse industrial settings.

Benchmarking and Comparative Studies: Conduct benchmarking and comparative studies to evaluate the performance of different cost function approaches and optimisation algorithms. Comparative studies can provide valuable insights into the strengths and limitations of various optimisation strategies and guide future research directions.

Industry-University Partnerships: Foster partnerships between academia and industry to bridge the gap between theoretical research and practical implementation. Collaborative projects can accelerate the adoption of cost function-driven sensor optimisation in real-world engineering applications.

At the end, these recommendations offer a roadmap for future research endeavours in sensor optimisation, emphasizing the continued advancement of cost function models, multi-objective optimisation techniques, and real-world deployment studies. By addressing these research areas, the field of sensor optimisation can further evolve, enabling engineers and decision-makers to make informed choices that balance system performance and cost-effectiveness in complex engineering systems.

2.8 References

[1] ‘Integrated Vehicle Health Management: Perspectives on an Emerging Field’. Accessed: May 03, 2023. [Online]. Available: https://www.sae.org/publications/books/content/r-405/

[2] A. Kulkarni, J. Terpenny, and V. Prabhu, ‘Sensor selection framework for designing fault diagnostics system’, *Sensors*, vol. 21, no. 19, Oct. 2021, doi: 10.3390/s21196470.

[3] K. Nakai, K. Yamada, T. Nagata, Y. Saito, and T. Nonomura, ‘Effect of Objective Function on Data-Driven Greedy Sparse Sensor Optimization’, *IEEE Access*, vol. 9, pp. 46731–46743, 2021, doi: 10.1109/ACCESS.2021.3067712.

[4] F. Wan, X. Yu, and Q. Yu, ‘Research on Optimal Sensor Placement for Aircraft Structural Health Management’, in *2017 IEEE International Conference on Prognostics and Health Management, ICPHM 2017*, Institute of Electrical and Electronics Engineers Inc., Jul. 2017, pp. 160–166. doi: 10.1109/ICPHM.2017.7998322.

[5] A. Krause, A. Singh, and C. Guestrin, ‘Near-Optimal Sensor Placements in Gaussian Processes: Theory, Efficient Algorithms and Empirical Studies’, 2008.

[6] E. Clark, T. Askham, S. L. Brunton, and J. Nathan Kutz, ‘Greedy Sensor Placement with Cost Constraints’, *IEEE Sens. J.*, vol. 19, no. 7, pp. 2642–2656, Apr. 2019, doi: 10.1109/JSEN.2018.2887044.

[7] T. Feng, R. Duan, Y. Lin, and Y. Zeng, ‘Optimal sensor placement based on multiattribute decision-making considering the common cause failure’, *Proc. Inst. Mech. Eng. C J. Mech. Eng. Sci.*, vol. 233, no. 19–20, pp. 7170–7182, Oct. 2019, doi: 10.1177/0954406219864982.

[8] C. Yang, R. Ma, and R. Ma, ‘Optimal Sensor Placement for Modal Identification in Multirotary-Joint Solar Power Satellite’, *IEEE Sens. J.*, vol. 20, no. 13, pp. 7337–7346, Jul. 2020, doi: 10.1109/JSEN.2020.2978081.

[9] Y. Wang, J. Xiang, Q. Mo, and S. He, ‘Compressed sparse time-frequency feature representation via compressive sensing and its applications in fault diagnosis’, *Measurement (Lond).*, vol. 68, pp. 70–81, 2015, doi: 10.1016/j.measurement.2015.02.046.

[10] T. L. Mohren, T. L. Daniel, S. L. Brunton, and B. W. Brunton, ‘Neural-inspired sensors enable sparse, efficient classification of spatiotemporal data’, *Proc. Natl. Acad. Sci. U. S. A.*, vol. 115, no. 42, pp. 10564–10569, Oct. 2018, doi: 10.1073/pnas.1808909115.

[11] P. Shi, X. Guo, D. Han, and R. Fu, ‘A sparse auto-encoder method based on compressed sensing and wavelet packet energy entropy for rolling bearing intelligent fault diagnosis’, *Journal of Mechanical Science and Technology*, vol. 34, no. 4, pp. 1445–1458, Apr. 2020, doi: 10.1007/s12206-020-0306-1.

[12] C. Jiang, Z. Chen, R. Su, and Y. C. Soh, ‘Group greedy method for sensor placement’, *IEEE Transactions on Signal Processing*, vol. 67, no. 9, pp. 2249–2262, May 2019, doi: 10.1109/TSP.2019.2903017.

[13] S. E. Otto and C. W. Rowley, ‘Inadequacy of Linear Methods for Minimal Sensor Placement and Feature Selection in Nonlinear Systems: A New Approach Using Secants’, *J. Nonlinear Sci.*, vol. 32, no. 5, Oct. 2022, doi: 10.1007/s00332-022-09806-9.

[14] W. Ostachowicz, R. Soman, and P. Malinowski, ‘Optimization of sensor placement for structural health monitoring: a review’, May 01, 2019, *SAGE Publications Ltd*. doi: 10.1177/1475921719825601.

[15] M. Pourali and A. Mosleh, ‘A functional sensor placement optimization method for power systems health monitoring’, *IEEE Trans. Ind. Appl.*, vol. 49, no. 4, pp. 1711–1719, 2013, doi: 10.1109/TIA.2013.2257978.

[16] M. M. Aral, J. Guan, and M. L. Maslia, ‘Optimal Design of Sensor Placement in Water Distribution Networks’, *J. Water Resour. Plan.*

Manag., vol. 136, no. 1, pp. 5–18, Jan. 2010, doi: 10.1061/(ASCE)WR.1943-5452.0000001.

[17] J. Y. Kim and J. H. Han, 'Optimal Transducer Placement for Deep Learning-Based Non-Destructive Evaluation', *Sensors*, vol. 23, no. 3, Feb. 2023, doi: 10.3390/S23031349.

[18] S.-M. Yang, J. Qiu, G.-J. Liu, and P. Yang, 'Sensor Selection and Optimization for Aerospace System Health Management under Uncertainty Testing'.

[19] L. M. Santi, T. S. Sowers, and R. B. Aguilar, 'Optimal sensor selection for health monitoring systems', in *41st AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit*, 2005. doi: 10.2514/6.2005-4485.

[20] W. A. Maul, G. Kopasakis, L. M. Santi, T. S. Sowers, and A. Chicatelli, 'Sensor selection and optimization for health assessment of aerospace systems', in *Collection of Technical Papers - 2007 AIAA InfoTech at Aerospace Conference*, American Institute of Aeronautics and Astronautics Inc., 2007, pp. 1443–1456. doi: 10.2514/6.2007-2849.

[21] T. S. Yoo and S. Lafortune, 'NP-completeness of sensor selection problems arising in partially observed discrete-event systems', *IEEE Trans. Automat. Contr.*, vol. 47, no. 9, pp. 1495–1499, 2002, doi: 10.1109/TAC.2002.802762.

[22] R. Fletcher, 'Practical Methods of Optimization', *Practical Methods of Optimization*, May 2000, doi: 10.1002/9781118723203.

[23] F. Guan, W. W. Cui, L. F. Li, and J. Wu, 'A comprehensive evaluation method of sensor selection for PHM based on grey clustering', *Sensors (Switzerland)*, vol. 20, no. 6, Mar. 2020, doi: 10.3390/s20061710.

[24] S. Joshi and S. Boyd, 'Sensor selection via convex optimization', *IEEE Transactions on Signal Processing*, vol. 57, no. 2, pp. 451–462, 2009, doi: 10.1109/TSP.2008.2007095.

[25] M. Shamaiah, S. Banerjee, and H. Vikalo, ‘Greedy sensor selection: Leveraging submodularity’, in *Proceedings of the IEEE Conference on Decision and Control*, Institute of Electrical and Electronics Engineers Inc., 2010, pp. 2572–2577. doi: 10.1109/CDC.2010.5717225.

[26] H. Wang, G. Pottie, K. Yao, and D. Estrin, ‘Entropy-based sensor selection heuristic for target localization’, *Third International Symposium on Information Processing in Sensor Networks, IPSN 2004*, pp. 36–45, 2004, doi: 10.1145/984622.984628.

[27] J. Xu, Y. Wang, and L. Xu, ‘PHM-Oriented Sensor Optimization Selection Based on Multiobjective Model for Aircraft Engines’, *IEEE Sens. J.*, vol. 15, no. 9, pp. 4836–4844, Sep. 2015, doi: 10.1109/JSEN.2015.2430361.

[28] N. Najjar, S. Gupta, J. Hare, S. Kandil, and R. Walthall, ‘Optimal sensor selection and fusion for heat exchanger fouling diagnosis in aerospace systems’, *IEEE Sens. J.*, vol. 16, no. 12, pp. 4866–4881, Jun. 2016, doi: 10.1109/JSEN.2016.2549860.

[29] X. Jiao, B. Jing, B. Jiao, S. Si, and Y. Wang, ‘An improved binary Wolf pack algorithm for solving optimal sensor selection problems’, in *Proceedings - 2018 Prognostics and System Health Management Conference, PHM-Chongqing 2018*, Institute of Electrical and Electronics Engineers Inc., Jan. 2019, pp. 162–170. doi: 10.1109/PHM-Chongqing.2018.00034.

[30] K. Manohar, J. N. Kutz, and S. L. Brunton, ‘Optimal Sensor and Actuator Selection Using Balanced Model Reduction’, *IEEE Trans. Automat. Contr.*, vol. 67, no. 4, pp. 2108–2115, Apr. 2022, doi: 10.1109/TAC.2021.3082502.

[31] Q. Yan and J. Chen, ‘Sensor selection via maximizing hybrid bayesian fisher information and mutual information in unreliable sensor networks’, *Electronics (Switzerland)*, vol. 9, no. 2, Feb. 2020, doi: 10.3390/ELECTRONICS9020283.

[32] R. Garnett, M. A. Osborne, and S. J. Roberts, ‘Bayesian optimization for sensor set selection’, *Proceedings of the 9th ACM/IEEE International Conference on Information Processing in Sensor Networks, IPSN ’10*, pp. 209–219, 2010, doi: 10.1145/1791212.1791238.

[33] M. F. R. Lee, K. Stanley, and Q. M. J. Wu, ‘Implementation of sensor selection and fusion using fuzzy logic’, *Annual Conference of the North American Fuzzy Information Processing Society - NAFIPS*, vol. 1, pp. 328–333, 2001, doi: 10.1109/nafips.2001.944273.

[34] Y. Wang, X. Song, and L. Zhang, ‘ISHM-oriented sensor optimization selection for aircraft engine system’, in *Advances in Intelligent Systems and Computing*, Springer Verlag, 2015, pp. 531–542. doi: 10.1007/978-3-662-47241-5_45.

[35] J. Xu, Y. Wang, and L. Xu, ‘PHM-Oriented Sensor Optimization Selection Based on Multiobjective Model for Aircraft Engines’, *IEEE Sens. J.*, vol. 15, no. 9, pp. 4836–4844, Sep. 2015, doi: 10.1109/JSEN.2015.2430361.

[36] Y. Hu, Z. Sun, L. Cao, Y. Zhang, and P. Pan, ‘Optimization configuration of gas path sensors using a hybrid method based on tabu search artificial bee colony and improved genetic algorithm in turbofan engine’, *Aerosp. Sci. Technol.*, vol. 112, p. 106642, May 2021, doi: 10.1016/J.AST.2021.106642.

[37] G. Jin and G. Jin, ‘Fault-diagnosis sensor selection for fuel cell stack systems combining an analytic hierarchy process with the technique order performance similarity ideal solution method’, *Symmetry (Basel).*, vol. 13, no. 12, Dec. 2021, doi: 10.3390/SYM13122366.

[38] P. Baraldi, G. Bonfanti, and E. Zio, ‘Differential evolution-based multi-objective optimization for the definition of a health indicator for fault diagnostics and prognostics’, *Mech. Syst. Signal Process.*, vol. 102, pp. 382–400, Mar. 2018, doi: 10.1016/j.ymssp.2017.09.013.

[39] M. A. Kassab, H. S. Taha, S. A. Shedied, and A. Maher, ‘A novel voting algorithm for redundant aircraft sensors’, *Proceedings of the World Congress on Intelligent Control and Automation (WCICA)*, vol. 2015-March, no. March, pp. 3741–3746, Mar. 2015, doi: 10.1109/WCICA.2014.7053339.

[40] W. Granig, D. Hammerschmidt, and H. Zangl, ‘Diagnostic coverage estimation method for optimization of redundant sensor systems’, *Proceedings of IEEE Sensors*, vol. 2017-December, pp. 1–3, Dec. 2017, doi: 10.1109/ICSENS.2017.8234088.

[41] W. Granig, L. M. Faller, D. Hammerschmidt, and H. Zangl, ‘Dependability considerations of redundant sensor systems’, *Reliab. Eng. Syst. Saf.*, vol. 190, Oct. 2019, doi: 10.1016/J.RESS.2019.106522.

[42] H. Zhou *et al.*, ‘Construction of health indicators for condition monitoring of rotating machinery: A review of the research’, Oct. 01, 2022, *Elsevier Ltd*. doi: 10.1016/j.eswa.2022.117297.

[43] H. J. Park, N. Ho Kim, and J.-H. Choi, *A Comparative Study of Health Monitoring Sensors based on Prognostic Performance*.

[44] M. A. Demetriou, ‘Economic aspects of sensor selection optimization of finite and infinite dimensional dynamical systems’, in *Proceedings of the American Control Conference*, Institute of Electrical and Electronics Engineers Inc., 2022, pp. 4139–4144. doi: 10.23919/ACC53348.2022.9867431.

[45] R. Moradi, S. Cofre-Martel, E. Lopez Droguett, M. Modarres, and K. M. Groth, ‘Integration of deep learning and Bayesian networks for condition and operation risk monitoring of complex engineering systems’, *Reliab. Eng. Syst. Saf.*, vol. 222, p. 108433, Jun. 2022, doi: 10.1016/J.RESS.2022.108433.

[46] Z. Gao, C. Cecati, and S. X. Ding, ‘A survey of fault diagnosis and fault-tolerant techniques-part I: Fault diagnosis with model-based and signal-

based approaches', Jun. 01, 2015, *Institute of Electrical and Electronics Engineers Inc.* doi: 10.1109/TIE.2015.2417501.

[47] G. Koutroulis, B. Mutlu, and R. Kern, 'Constructing robust health indicators from complex engineered systems via anticausal learning', *Eng. Appl. Artif. Intell.*, vol. 113, Aug. 2022, doi: 10.1016/j.engappai.2022.104926.

[48] S. L. Sun and J. Ma, 'Optimal filtering and smoothing for discrete-time stochastic singular systems', *Signal Processing*, vol. 87, no. 1, pp. 189–201, Jan. 2007, doi: 10.1016/J.SIGPRO.2006.05.007.

[49] S. Ding, G. Tian, V. Moorthy, and P. Wang, 'New feature extraction for applied stress detection on ferromagnetic material using magnetic Barkhausen noise', doi: 10.1016/j.measurement.2015.05.031.

[50] P. D. Welch, 'The Use of Fast Fourier Transform for the Estimation of Power Spectra: A Method Based on Time Averaging Over Short, Modified Periodograms', *IEEE Transactions on Audio and Electroacoustics*, vol. 15, no. 2, pp. 70–73, 1967, doi: 10.1109/TAU.1967.1161901.

[51] M. Gerdes, D. Galar, and D. Scholz, 'Decision trees and the effects of feature extraction parameters for robust sensor network design', *Eksploatacja i Niezawodnosc*, vol. 19, no. 1, pp. 31–42, 2017, doi: 10.17531/EIN.2017.1.5.

[52] A. Widodo and B. S. Yang, 'Application of nonlinear feature extraction and support vector machines for fault diagnosis of induction motors', *Expert Syst. Appl.*, vol. 33, no. 1, pp. 241–250, Jul. 2007, doi: 10.1016/J.ESWA.2006.04.020.

[53] O. O. Akinola, A. E. Ezugwu, J. O. Agushaka, R. A. Zitar, and L. Abualigah, 'Multiclass feature selection with metaheuristic optimization algorithms: a review', Nov. 01, 2022, *Springer Science and Business Media Deutschland GmbH*. doi: 10.1007/s00521-022-07705-4.

[54] G. K. K. Lee, H. Kasim, R. P. Sirigina, S. S. Q. How, S. King, and T. G. G. Hung, ‘Smart Robust Feature Selection (SoFt) for imbalanced and heterogeneous data’, *Knowl. Based. Syst.*, vol. 236, Jan. 2022, doi: 10.1016/j.knosys.2021.107197.

[55] H. Peng, F. Long, and C. Ding, ‘Feature selection based on mutual information: Criteria of Max-Dependency, Max-Relevance, and Min-Redundancy’, *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 27, no. 8, pp. 1226–1238, Aug. 2005, doi: 10.1109/TPAMI.2005.159.

[56] B. Rezaeianjouybari and Y. Shang, ‘Deep learning for prognostics and health management: State of the art, challenges, and opportunities’, *Measurement (Lond).*, vol. 163, Oct. 2020, doi: 10.1016/j.measurement.2020.107929.

[57] O. O. Aremu, D. Hyland-Wood, and P. R. McAree, ‘A machine learning approach to circumventing the curse of dimensionality in discontinuous time series machine data’, *Reliab. Eng. Syst. Saf.*, vol. 195, Mar. 2020, doi: 10.1016/J.RESS.2019.106706.

[58] N. M. Thoppil, V. Vasu, and C. S. P. Rao, ‘Deep Learning Algorithms for Machinery Health Prognostics Using Time-Series Data: A Review’, *Journal of Vibration Engineering & Technologies 2021 9:6*, vol. 9, no. 6, pp. 1123–1145, Mar. 2021, doi: 10.1007/S42417-021-00286-X.

[59] S. R. Saufi, Z. A. Bin Ahmad, M. S. Leong, and M. H. Lim, ‘Challenges and opportunities of deep learning models for machinery fault detection and diagnosis: A review’, *IEEE Access*, vol. 7, pp. 122644–122662, 2019, doi: 10.1109/ACCESS.2019.2938227.

[60] Y. Lei, B. Yang, X. Jiang, F. Jia, N. Li, and A. K. Nandi, ‘Applications of machine learning to machine fault diagnosis: A review and roadmap’, Apr. 01, 2020, *Academic Press*. doi: 10.1016/j.ymssp.2019.106587.

[61] Q. Yang, Y. Zhang, W. Dai, and S. J. Pan, ‘Transfer Learning’, *Transfer Learning*, Jan. 2020, doi: 10.1017/9781139061773.

[62] V. Zaccaria, M. Rahman, I. Aslanidou, and K. Kyprianidis, ‘A review of information fusion methods for gas turbine diagnostics’, *Sustainability (Switzerland)*, vol. 11, no. 22, Nov. 2019, doi: 10.3390/SU11226202.

[63] W. Yan, K. Goebel, and K. F. Goebel, ‘Sensor validation and fusion for gas turbine vibration monitoring’, *https://doi.org/10.1117/12.487206*, vol. 5107, pp. 106–117, Aug. 2003, doi: 10.1117/12.487206.

[64] A. J. Volponi Pratt, E. Hartford, C. Tom Brotherton, D. L. Simon, and U. Army, ‘Development of an Information Fusion System for Engine Diagnostics and Health Management’, 2004. Accessed: May 24, 2023. [Online]. Available: http://www.sti.nasa.gov

[65] A. Fritze, U. Mönks, C. A. Holst, and V. Lohweg, ‘An approach to automated fusion system design and adaptation’, *Sensors (Switzerland)*, vol. 17, no. 3, Mar. 2017, doi: 10.3390/s17030601.

3 Normalised Diagnostic Contribution Index (NDCI) Integration to Multi Objective Sensor Optimisation Framework (MOSOF) – ECS Case

Chapter 3 (based on the second research paper) introduces and applies the Normalised Diagnostic Contribution Index (NDCI) within a Multi-Objective Sensor Optimisation Framework (MOSOF) to an aircraft Environmental Control System (ECS) case study. This chapter builds on the literature insights by proposing a domain-specific sensor ranking metric (NDCI) that captures each sensor's fault diagnostic value – combining its fault separation power, sensitivity to fault severity, and uniqueness (non-redundancy) – into a single score. The NDCI is then integrated into a multi-objective genetic algorithm framework to optimise sensor selection while considering real-world constraints and stakeholder criteria. In a high-fidelity Boeing 737-800 ECS simulation (covering multiple fault modes and severities), the NDCI-driven optimisation is compared against a baseline mRMR approach. The results demonstrate that NDCI can achieve equivalent or better diagnostic performance with a more compact sensor suite. Notably, the top NDCI-based solution required only three sensors to achieve high fault coverage, whereas mRMR's recommendation needed six sensors for a similar performance level. This indicates NDCI's emphasis on diagnostic value over simple redundancy, expanding the feasible solution space to smaller, more informative sensor sets. The chapter also considers multiple stakeholder perspectives (OEM, airline, and MRO), demonstrating that the enhanced NDCI–MOSOF framework can accommodate different priorities, such as cost, reliability, and benefit-to-cost ratios. Overall, Chapter 3 validates the efficacy of the NDCI approach in an aerospace subsystem, improving diagnostics efficiency and demonstrating a scalable, multi-stakeholder sensor optimisation strategy. This lays the groundwork for extending the methodology to a complete aircraft platform with interacting subsystems, which is explored in the next chapter.

3.1 Introduction

The Modern aerospace and power generation systems increasingly depend on sophisticated sensor networks to ensure reliable fault detection, rapid isolation, and timely maintenance actions [1]. Advances in sensor technology have enabled the deployment of extensive sensor suites across a wide array of applications—from industrial process control and environmental monitoring to security and safety-critical operations [1, 2]. In these complex systems, overall performance hinges on the careful selection and configuration of sensors—a process that must balance conflicting objectives such as diagnostic accuracy, cost efficiency, reliability, and integration complexity. For instance, while original equipment manufacturers (OEMs) emphasise design simplicity and certification compliance, airline operators focus on minimising operational disruptions and maintenance expenses, and Maintenance, Repair, and Overhaul (MRO) providers prioritise swift fault isolation [2].

Figure 3-1 illustrates the diverse domains where sensors play a pivotal role, including diagnostics, control, monitoring, safety, and system optimisation. Diagnostic applications benefit from robust fault detection and predictive maintenance schedules; control systems require real-time adjustments to optimise performance; monitoring tasks ensure environmental compliance and accurate data logging; and safety systems are essential for protecting assets and personnel.

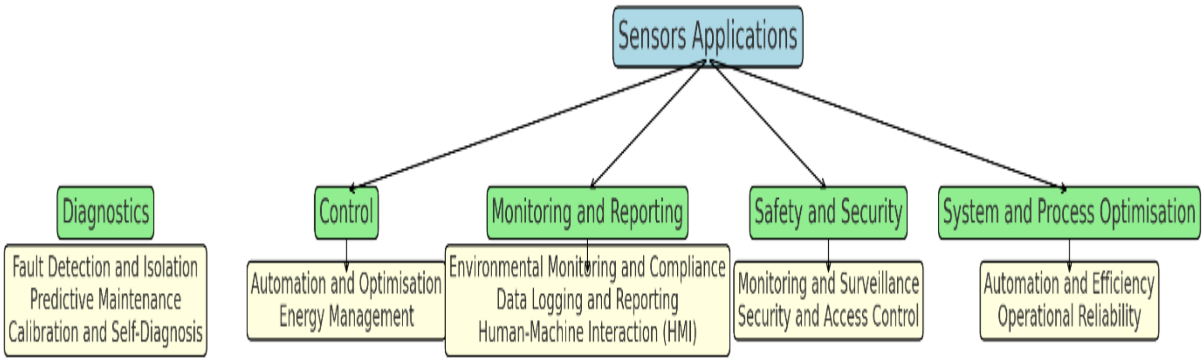

Figure 3-1 Sensor Application Domains

Notably, within Figure 3-1, the 'Diagnostic' block is intentionally depicted as structurally decoupled from the generic sensor data streams. This visual distinction underscores a core premise of this research: optimal sensor selection cannot be driven by a universal, domain-agnostic approach. While each application domain highlighted in the figure theoretically requires its own bespoke evaluation metric to effectively guide sensor optimisation toward its unique operational goals, this thesis isolates and focuses exclusively on the diagnostic domain. The unconnected block thus signifies the necessity of pulling diagnostics out of the generic sensor framework to develop a dedicated, domain-specific metric (such as the NDCI) that optimises sensor network architectures strictly for guiding its application, in this case, fault detection, severity tracking, and isolation.

Traditional sensor selection strategies have typically focused on either maximising diagnostic accuracy through high redundancy or reducing sensor counts to minimise costs. However, approaches that concentrate solely on one objective often prove suboptimal. High-redundancy configurations can drive up expenses without corresponding improvements in fault detection, while overly cost-driven designs risk undermining coverage of critical components [2]. Techniques such as the Minimum Redundancy Maximum Relevance (mRMR) algorithm have been widely employed to balance high relevance with low redundancy, thereby identifying sensor signals that are both informative and non-overlapping [3]. Recent enhancements to mRMR, including energy optimisation for wireless sensor networks [4] and improved measures of redundancy and relevance [5], further refine its performance in feature selection and classification tasks. Furthermore, several design optimisation frameworks have been proposed to integrate testability and maintainability requirements into sensor selection processes, as illustrated by works on optimal sensor selection for health monitoring systems [6] and aerospace system health management under uncertainty [7].

Recent investigations into aircraft Environmental Control Systems (ECS) have highlighted these challenges. Legacy ECS designs—exemplified by early

Boeing 737 models that utilised only a few temperature sensors for control [8]—often struggle to detect subtle faults, such as heat exchanger fouling or valve leakage. Empirical studies have demonstrated that sensor selection methods based solely on traditional criteria may underperform in the presence of multiple fault modes and varying operational conditions, as shown in cross-condition fault diagnosis studies [9]. In response, researchers have developed advanced multi-objective optimisation techniques that integrate comprehensive diagnostic metrics with practical constraints. For example, integrated combinatorial approaches employing reduced-order models and Gaussian processes have been proposed to further enhance sensor placement optimisation [10].

Central to these advanced methodologies is the Normalised Diagnostic Contribution Index (NDCI), which consolidates critical attributes—such as sensitivity, fault-separation capability, and uniqueness—into a single quantitative measure of diagnostic performance [11, 12]. When integrated within the Multi-Objective Sensor Optimisation Framework (MOSOF) [13], the NDCI enables a holistic evaluation of each sensor's contribution to fault detection, thereby facilitating a balanced assessment that accounts for both technical performance and resource constraints. In parallel, PHM-oriented sensor optimisation models, such as those applied to aircraft engines, have shown promising results in aligning sensor selection with prognostics and health management objectives [14]. Advanced fault diagnosis techniques—like subspace approaches to multidimensional fault identification and reconstruction—complement these frameworks by enhancing the capability to detect and isolate faults across multiple dimensions [15].

Building on these developments, this work embeds the NDCI within MOSOF and employs a Multi-Objective Genetic Algorithm (MOGA) to refine the sensor selection process for complex diagnostic applications. This enhanced framework simultaneously optimises four critical objectives:

1- Overall Diagnostic Performance (Maximising Observability): This objective acts as the primary technical driver, evaluated via a composite index anchored by the Normalised Diagnostic Contribution Index (NDCI). Rather than simply counting detectable faults, this objective quantifies the severity-sensitive fault-separation capability and information uniqueness of the selected sensor suite, ensuring maximum diagnostic yield with minimal data redundancy.

2- Lifecycle Cost of Sensor Deployment (Minimising Economic Penalty): This objective transcends initial unit acquisition costs to account for the broader economic integration penalties. It specifically targets the minimisation of both the direct financial expenses of hardware procurement and the indirect, long-term operational costs associated with parasitic weight (e.g., cumulative fuel burn penalties induced by additional sensors and wiring harnesses).
3- Reliability and Operational Efficiency (Maximising Suite MTBF): Adding sensors inherently increases the mathematical probability of a sensor failure. Therefore, this objective focuses on mitigating parasitic unreliability by maximising the aggregate MTBF of the sensor network. It ensures that the diagnostic suite enhances system robustness rather than becoming a primary source of false alarms or maintenance burdens.
4- Stakeholder-Specific Secondary Criteria (Optimising Contextual Constraints): This objective captures the diverging operational priorities of the aerospace ecosystem. For OEMs, it evaluates architectural compatibility and certifiability risks. For Airlines and MROs, it optimises for downstream benefit-to-cost ratios, fleet availability, and the sensor suite's ability to facilitate rapid fault isolation to reduce aircraft turnaround times [16].

In practice, while traditional methods like mRMR may suggest larger sensor sets to guarantee redundancy, the NDCI-based approach has demonstrated the ability to achieve comparable fault detection with fewer sensors, thereby enhancing overall diagnostic efficiency. The optimisation process addresses the NP-complete nature of sensor placement [13] by initially employing mRMR as a pre-processing step to eliminate redundant signals and subsequently applying MOGA to navigate the complex combinatorial search space, thereby deriving Pareto-optimal solutions that balance multiple objectives [17]. System-level fault diagnosis studies applied to aircraft ECS further reinforce the value of such integrated optimisation approaches [18]. Moreover, multi-objective sensor placement strategies have been explored for composite aircraft structures using Kriging-based approaches, which provide valuable insights into optimal sensor layouts in complex material systems [19].

Case studies further validate the efficacy of this integrated approach. For instance, simulation studies on a Boeing 737-800 ECS have demonstrated that augmenting legacy configurations with strategically selected sensors, guided by the NDCI, can substantially improve fault detection and isolation capabilities [20]. Similarly, multi-objective optimisation strategies have been successfully applied to turbofan engine monitoring, where an optimised sensor set reduced 35 variables to just four critical indicators without compromising diagnostic accuracy [21]. Other domains, such as helicopter drive systems and polymer electrolyte membrane fuel cell prognostics, have also benefited from these methodologies [22, 23]. Additionally, innovations like Teledyne's Aircraft Cabin Environment Sensor (ACES) for the Boeing 737 exemplify the trend toward comprehensive sensor optimisation, balancing extensive environmental monitoring with practical constraints [20]. Design-for-diagnosability practices in modern systems, such as those implemented in the Boeing 787's electric ECS, further underscore the importance of rigorous multi-objective trade studies in justifying sensor additions with respect to weight and power consumption [24, 25].

Moreover, intellectual property considerations highlight the competitive advantages of optimised sensor layouts, with several patents underscoring the value of these innovations [26, 27]. In summary, the evolution from traditional, single-objective sensor selection to sophisticated multi-objective frameworks, such as MOSOF, represents a significant advancement in the design and operation of complex engineering systems. By integrating diagnostic metrics, such as the NDCI, with multi-objective optimisation algorithms—and by leveraging emerging techniques in sensor network design and fault diagnosis—the approach not only expands the feasible solution space for sensor configurations but also meets the diverse requirements of OEMs, airlines, and MRO stakeholders. This work enhances fault detection and system reliability while laying the foundation for future research aimed at refining multi-stakeholder decision-making in sensor network design.

3.2 NDCI Integration into MOSOF

3.2.1 Rationale for NDCI-Centric Sensor Evaluation

T MOSOF was developed as a general approach to sensor implementation in complex systems, addressing multiple performance, cost, and reliability criteria.

However, MOSOF’s broad scope sometimes masks the specific performance nuances essential for applications requiring precise sensor differentiation. The NDCI has emerged due to this work as a robust metric that quantifies sensor performance by evaluating parameters such as signal separation, sensitivity, and uniqueness. Integrating NDCI into MOSOF refines the sensor evaluation process by introducing an objective, quantitative measure of sensor quality that is particularly beneficial in applications where subtle performance differences can significantly impact system outcomes.

Incorporating NDCI into MOSOF is especially pertinent for modern applications, ranging from automotive systems to industrial monitoring, where the relative performance of sensors under fault conditions is crucial. By using NDCI as the core performance metric, the enhanced framework becomes more sensitive to changes in sensor behaviour, thereby improving sensor selection and configuration decisions. This integration addresses the limitations of a one-size-fits-all approach by tailoring the framework to meet the specific requirements of diverse diagnostic applications.

Figure 3-2 presents the detailed thermodynamic block diagram of the Boeing 737-800 ECS PACK, developed within Cranfield University's System Environment for SESAC simulation. The schematic illustrates the pneumatic pathway where high-pressure engine bleed air enters the system via the primary valve (PV), regulated by a dedicated controller to maintain a target mass flow of 0.4556 kg/s. The airflow is subsequently divided into two distinct paths: a primary cooling stream that is processed through the Primary Heat Exchanger (PHX), the Air Cycle Machine (ACM), and the Secondary Heat Exchanger (SHX), and a hot bypass stream regulated by the Temperature Control Valve (TCV). The cold stream also traverses the High-Pressure Water Separator (HPWS) sub-loop—comprising the Reheater (RHX), Condenser (CHX), and a water separator—to extract moisture before mixing. Cooling for the PHX and SHX is provided by a crossflow of external ram air. Crucially for diagnostic evaluation, the architecture embeds explicit fault injection nodes (labelled 'Faulty Ram', 'Faulty PHX', 'Faulty SHX', and 'ACM'), which allow for the simulated introduction of progressive physical degradation, such as heat exchanger fouling or reduced ACM mechanical efficiency. Ultimately, the conditioned cold air and the hot bypass air are combined in the merge module. A virtual temperature sensor monitors this mixed output and provides feedback to the controller, which continuously modulates the TCV to achieve the required target supply temperature of 291 K at a nominal cabin pressure of 79,500 Pa [28]. This high-fidelity digital twin provides the foundational, severity-dependent sensor data required for the subsequent NDCI calculations.

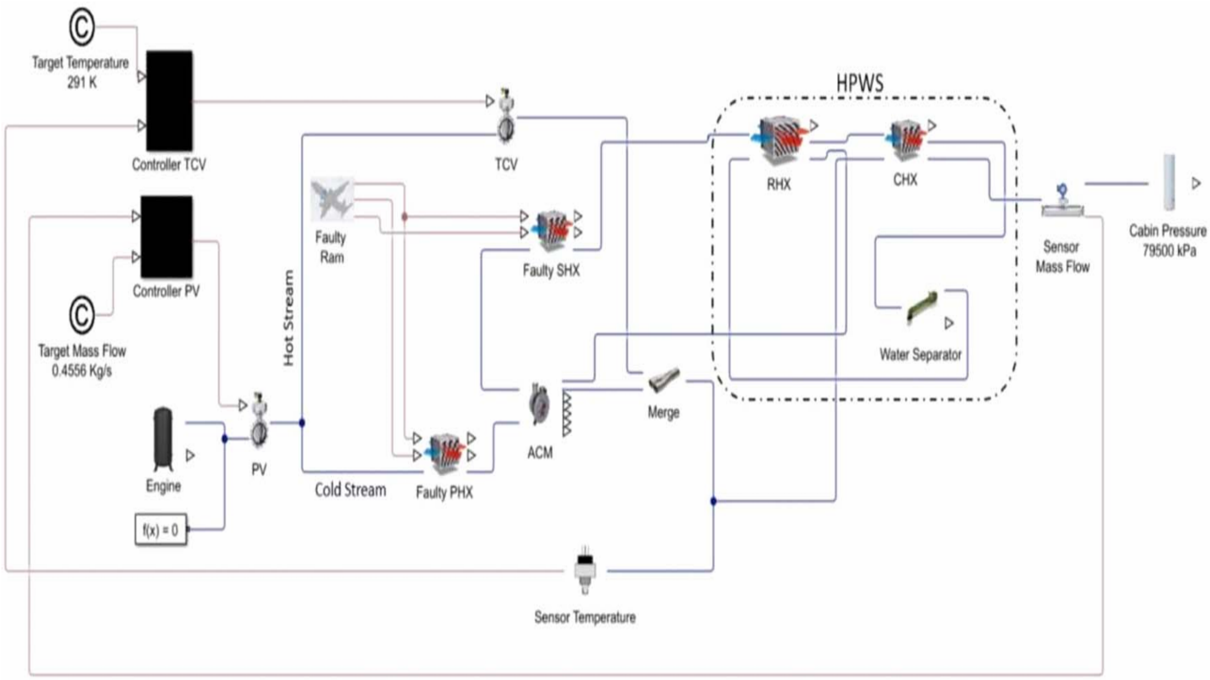

Figure 3-2 SESAC fault simulation model for the B737-800 PACK [28]

Figure 3-3 shows temperature sensor readings across various ECS components under different fault scenarios—PHX, SHX, ACM, and Ram Air—derived from the SESAC simulation which is developed in Cranfield University IVHM department [28].

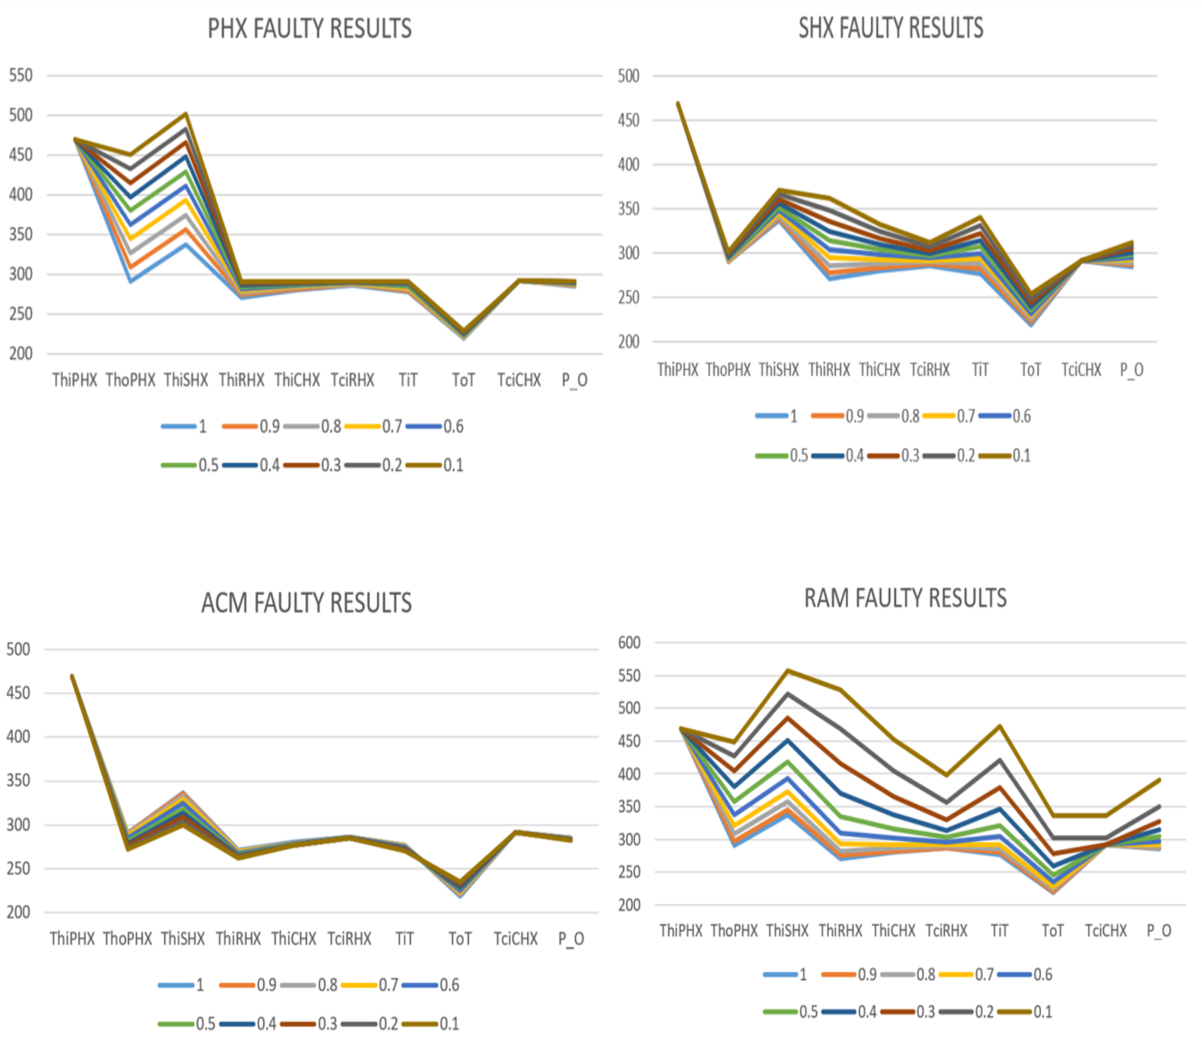

Figure 3-3 Temperature sensors' readings across ECS components obtained from the SESAC simulation for four fault modes [28]

The line plots in the Figure 3-3 demonstrate how sensor outputs shift relative to healthy baseline values as fault severity scales from 0.1 to 1.0. Specifically, these visualisations reveal that parameters such as the Hot Inlet Temperature to the Secondary Heat Exchanger (ThiSHX) and the Turbine Outlet Temperature (ToT) consistently exhibit notable deviations and wide proportional spreads across multiple fault modes, signalling high diagnostic relevance. Conversely, parameters such as the Cold Inlet Temperature (TciCHX) remain largely static, indicating poor observability for these specific degradation paths. This distinct variation in signal spread forms the empirical groundwork for the subsequent NDCI-based evaluations.

Furthermore, it is critical to acknowledge the fundamental reliance of this methodology on the availability of progressive degradation data. If continuous degradation profiles (or explicit severity labels) were not available upfront, the Sensitivity (S) component of the NDCI framework could not be mathematically derived. In the absence of severity data, the evaluation would be forced to collapse into a binary fault/no-fault classification problem, relying solely on baseline Separation Power or standard mutual information algorithms (e.g., mRMR). Consequently, the optimisation results would fundamentally shift: the framework might successfully select sensors capable of triggering an initial fault alarm, but it would systematically fail to prioritize sensors with the high-resolution tracking capabilities required for progressive prognostics and Remaining Useful Life (RUL) estimation.

3.2.2 Theoretical Underpinnings of NDCI

NDCI is computed by combining several key metrics that collectively capture the sensor's ability to distinguish between normal and faulty operating conditions. Typically, the index is derived from three components:

Separation Power (SP): For each sensor, SP is calculated by taking the absolute difference between the sensor's output under fault conditions and its healthy baseline, then normalising this difference by the range of the healthy values. Mathematically, if [yf] is the sensor reading under fault and [yh] is the healthy reading, SP is given by:

$$\text{SP} \ \frac{|y_f - y_h\,|}{range(y_h) + \epsilon} \qquad \textbf{(3-1)}$$

Where ϵ\epsilonϵ is a small constant to avoid division by zero.

Sensitivity (S): Sensitivity measures the relative change in sensor output in response to deviations in system performance. This is computed by normalising the same absolute difference by the complement of a fault severity indicator (or degradation level), ensuring that even small but significant changes are recognised:

$$S = \frac{|y_f - y_h|}{(1-Deg) + \epsilon} \quad \textbf{(3-2)}$$

where Deg represents the normalised degradation level.

It is critical to acknowledge the operational implications of this formulation. The inclusion of the degradation parameter assumes an a priori understanding of fault severity, typically derived from high-fidelity simulations or richly labelled historical datasets. In real-world legacy systems where continuous fault severity indicators are unavailable, this mathematical framework is fundamentally restricted. Without an understanding of progressive fault severity, the Sensitivity component degrades into a simple absolute residual, blinding the NDCI metric to prognostic tracking capabilities. Consequently, the optimisation algorithm would inherently bias towards sensors that react strongly to the binary presence of a fault, rather than those possessing the granular resolution required to track subtle, monotonic degradation over time. This highlights the vital dependency between high-quality fault characterisation data and optimal sensor network design.

Uniqueness (U): Uniqueness assesses the degree to which a sensor's responses differ from those of other sensors. By calculating the average Euclidean distance between the sensor's fault-response vector and those of its peers, and then normalising by the maximum observed distance, U is defined as:

$$U = \frac{avgDistance}{\max{(avgDistance)}} \quad \textbf{(3-3)}$$

Mathematically, NDCI is typically formulated as the average of these components:

$$(NDCI = (SP + S + U) / 3) \quad \textbf{(3-4)}$$

Although alternative weighted formulations can be adopted depending on the application requirements, this composite metric thereby captures the multifaceted performance characteristics of a sensor in a single scalar value.

Figure 3-4 presents the computed SP, S, and U values for the four ECS fault modes (PHX, SHX, ACM, and RAM), alongside the resulting NDCI for each sensor. The bar charts highlight which sensors excel in fault separation and sensitivity, as well as how unique each sensor's response is relative to others. This information is crucial for identifying sensor candidates with high diagnostic relevance.

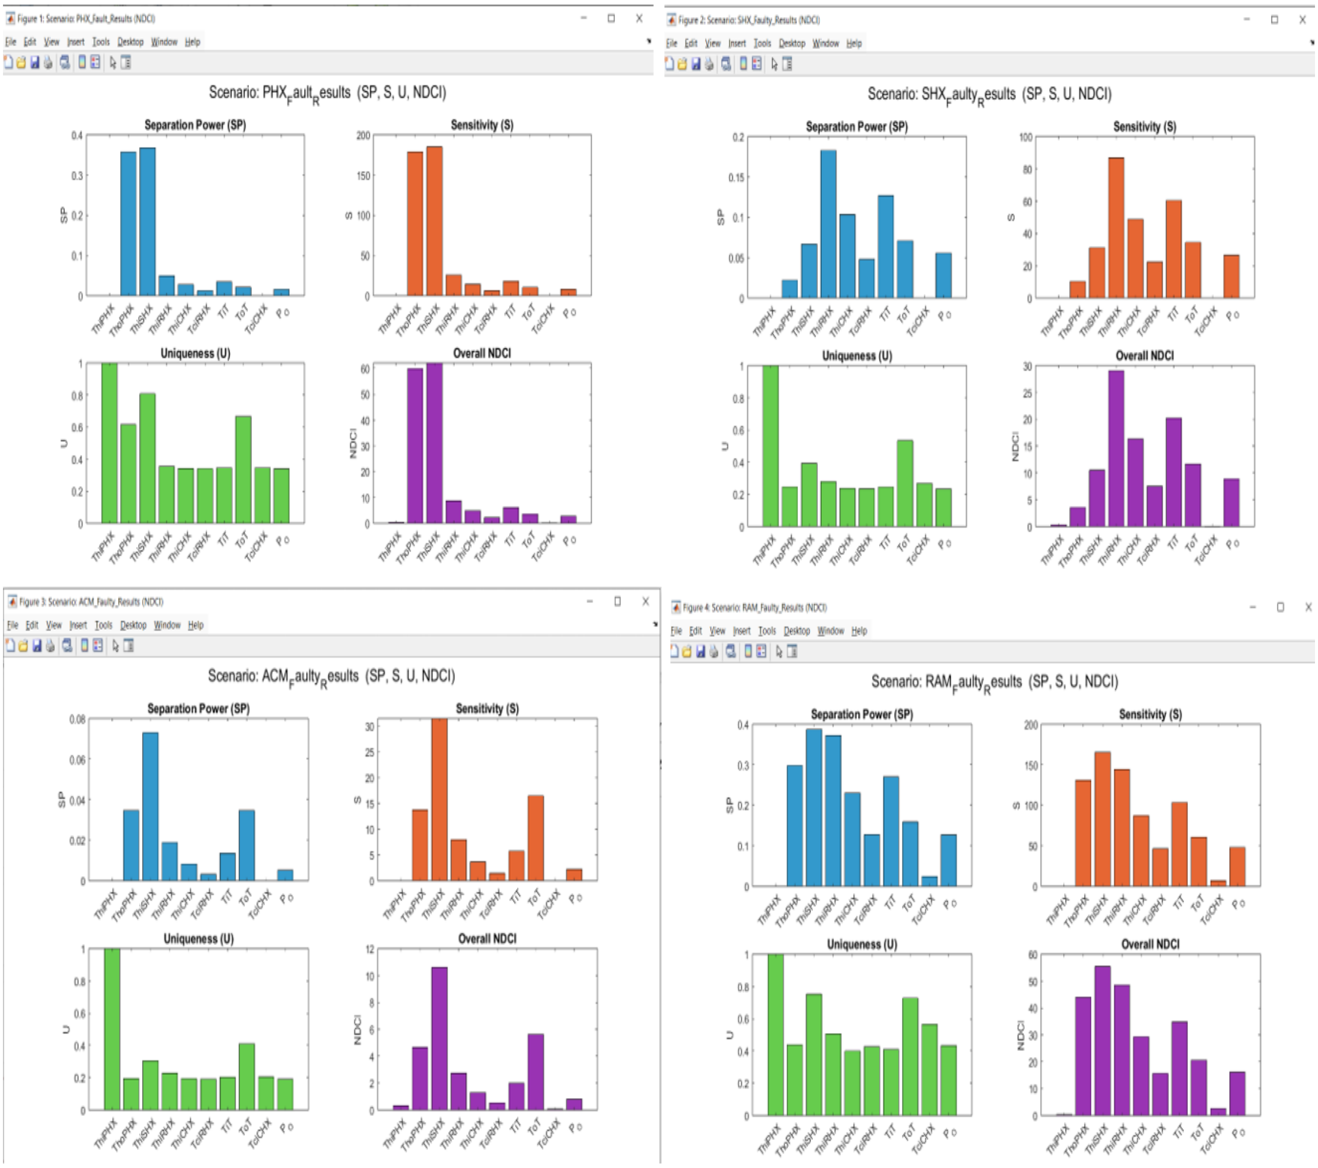

Figure 3-4 NDCI calculations for the four ECS fault modes

Figure 3-4 explicitly delineates the constituent components of the NDCI across four distinct ECS fault scenarios: PHX, SHX, ACM, and Ram Air faults. For each fault mode, the figure provides a quadrant of bar charts mapping the individual diagnostic metrics, SP, S, and U, against the candidate sensor suite plotted on the x-axis.

These visualisations deconstruct the underlying mechanics of the NDCI algorithm. The SP and S subplots demonstrate a given sensor's absolute deviation from the baseline and its severity-tracking resolution, respectively, while the U subplot quantifies the orthogonality of its information relative to the rest of the network. The final 'Overall NDCI' subplot mathematically synthesises these three parameters into a single, aggregated diagnostic utility score for each sensor. By presenting these metrics side-by-side, the figure visually confirms how sensors with high raw sensitivity might still receive a moderated overall NDCI score if their data is highly redundant (low uniqueness). This comprehensive breakdown ensures that the subsequent MOSOF selection phase prioritises a diverse, diagnostically robust sensor topology rather than merely selecting the most reactive individual transducers.

Figure 3-5 compares NDCI-based rankings against mRMR-based rankings across the same fault scenarios. The side-by-side bar charts illustrate instances where certain sensors score well in both metrics, indicating strong overall diagnostic utility, and cases where a sensor’s performance may be more specialised - contributing unique information that a single selection criterion could overlook. These comparisons underscore the importance of a multi-criteria approach to sensor evaluation.

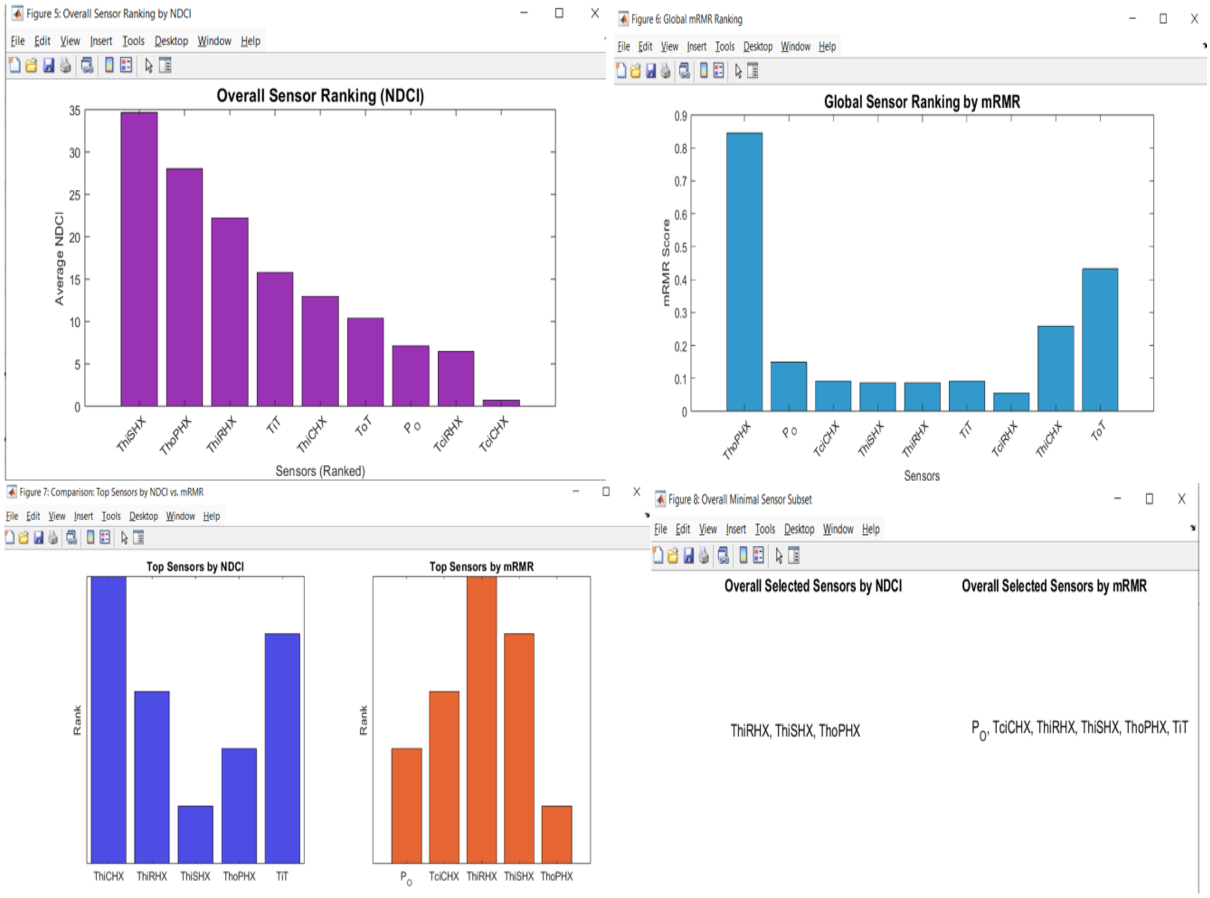

Figure 3-5 NDCI and mRMR score comparison across fault modes

NDCI and mRMR results derived from the SESAC simulation are shown in Figures 3-4 and 3-5. MATLAB code for these simulations is available on GitHub [29], facilitating reproducibility and further exploration of the data.

3.2.3 Proposed Methodology for NDCI-MOSOF Integration

The integration of NDCI into MOSOF involves embedding the NDCI within the multi-objective optimisation formulation. In the enhanced framework, NDCI is used as the central diagnostic performance index in the first objective function, which is combined with additional objectives such as cost, reliability (e.g., MTBF or efficiency), and secondary factors (e.g., compatibility, coverage, or benefit-to-cost ratio). For instance, in an OEM scenario, the performance objective is defined as a weighted combination of cost, NDCI and sensor accuracy. In contrast, in Airline and MRO scenarios, the performance objective combines

NDCI with coverage metrics, while additional objectives capture reliability and efficiency.

The flowchart in Figure 3-6 illustrates the updated procedure for computing NDCI and selecting optimal sensor sets under multiple fault scenarios in the MATLAB environment. Compared to the earlier version, this refined approach encompasses a multi-scenario perspective, integrates a global mRMR ranking, and incorporates coverage-based subset selection. The following steps summarise the methodology:

Input Data Preparation:

The process begins by loading multi-scenario data from an Excel file (e.g., *SESAC1.xlsx*). Each sheet corresponds to a distinct scenario (or fault type), including a healthy baseline row and multiple fault rows. This step consolidates all necessary information—sensor readings, fault severity indicators (deg), and baseline values—into a structured format.

Scenario Processing (SP, S, U Computation):

For each scenario, the algorithm calculates three metrics for every sensor:

Separation Power (SP): Normalised difference between fault and healthy measurements.

Sensitivity (S): Relative change in sensor output scaled by the fault severity.

Uniqueness (U): Distinctiveness of a sensor's fault signatures compared to those of other sensors in the same scenario.

These metrics are then averaged to produce the scenario-level NDCI for each sensor, providing a quantitative measure of diagnostic performance across multiple fault conditions.

Filtering Constant Sensors:

After processing all scenarios, sensors that exhibit negligible variation (below a specified threshold) are removed to eliminate redundant or uninformative signals. This ensures that subsequent analyses focus only on sensors capable of contributing meaningful diagnostic information.

Aggregating Scenario Data for Global Analysis:

The NDCI values and the underlying fault data from each scenario are compiled into global matrices. This consolidated view facilitates system-wide comparisons and ensures consistency when deriving overall sensor rankings.

Overall NDCI Ranking:

An overall average NDCI is computed across all scenarios, and sensors are ranked accordingly. This ranking highlights those sensors that consistently demonstrate high diagnostic value in multiple fault conditions.

Coverage-Based Subset Selection (Per Scenario):

Each scenario undergoes a coverage analysis, which measures how effectively a subset of sensors (chosen in descending order of their NDCI or mRMR scores) captures the scenario's total diagnostic capability. Once the subset meets a predefined coverage threshold (e.g., 95% of the maximum possible NDCI), no further sensors are added.

Union of Minimal Subsets:

The minimal subsets identified per scenario are combined to form a single, system-wide sensor set. This ensures that unique fault modes in different scenarios are adequately covered without unnecessarily inflating sensor counts.

Final Sensor Selection and Output:

The resulting sensor set, along with the intermediate rankings (NDCI-based and mRMR-based) and coverage analyses, is presented as the final output. This includes visualisations comparing top-k sensors by NDCI versus mRMR and any additional metrics relevant to cost or reliability.

By adopting this flow, the approach ensures robust sensor selection across multiple fault scenarios, balances redundancy with coverage, and incorporates NDCI insights. The methodology thus provides a comprehensive framework that can be adapted to diverse domains, ranging from aerospace to industrial process monitoring.

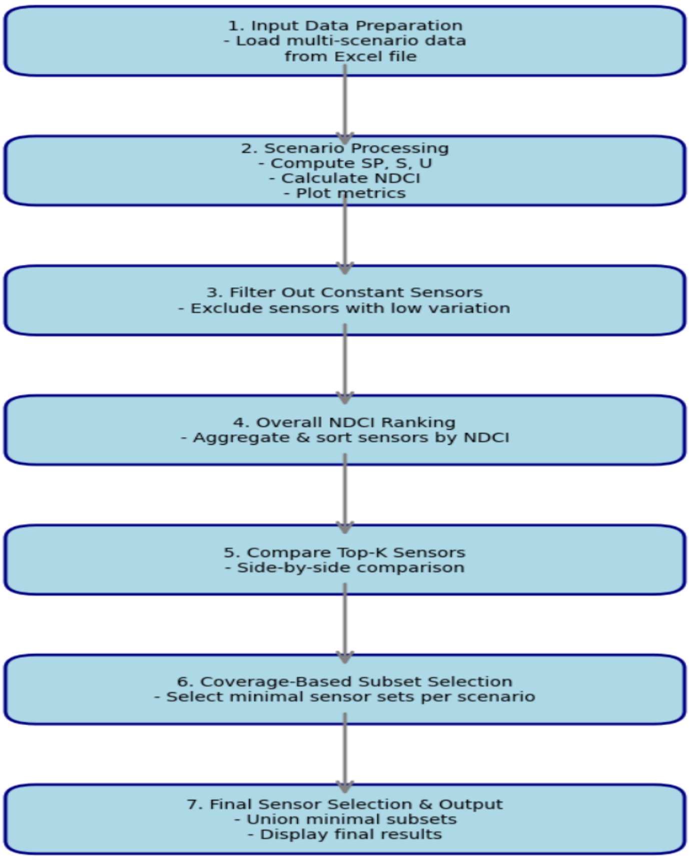

Figure 3-6 NDCI Flowchart

This methodology is operationalised via an MOGA that selects sensor pairs meeting specific constraints (e.g., budget limits and minimum performance thresholds). The optimisation process generates a Pareto front that represents trade-offs among the objectives. Detailed simulation studies demonstrate that this NDCI-MOSOF integration not only broadens the feasible solution space but also yields sensor configurations that are more attuned to application-specific performance needs. Research methodology is presented in Figure 3-7.

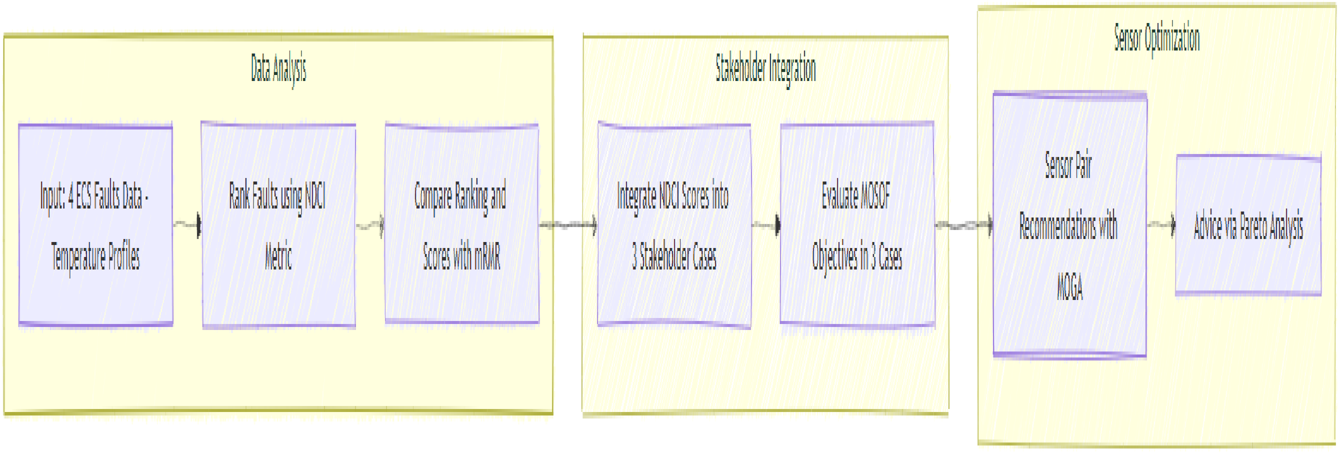

Figure 3-7 Research Methodology

Grounding sensor evaluation in NDCI enables a more precise and application-oriented sensor selection process, thus paving the way for future research in domain-specific sensor optimisation strategies.

3.3 Multi-Objective Optimisation Framework

The sensor selection problem is formulated as a multi-objective optimisation task where a binary decision vector represents the inclusion or exclusion of each sensor in the network. In the formulation, each sensor in the pool is associated with a set of performance attributes—most notably, the NDCI, along with secondary criteria such as accuracy, cost, reliability (or efficiency), and compatibility (or coverage). The objective is to identify sensor pairs that yield optimal trade-offs among these metrics while satisfying system-level constraints.

Table 3-1 offers a focused, NDCI-centric perspective on how the three primary aviation stakeholders, OEM, Airlines, and MRO, integrate diagnostic performance into their operational objectives. By placing NDCI at the forefront of each decision-making criterion, the table captures both the high-level trade-offs and specific considerations (e.g., cost, reliability, and downtime) that influence sensor optimisation strategies. This view not only underscores the importance of aligning sensor choices with each stakeholder’s priorities but also clarifies how a unified performance metric, such as NDCI, can streamline communication and goal setting across the aircraft's entire lifecycle.

Main Application Areas	Key Functions	Examples	MOSOF: Sensor Selection	MOSOF: Sensor Placement	MOSOF: Data Processing	MOSOF: Sensor Operation
Diagnostics	Fault detection, predictive maintenance, calibration	Detecting machine wear, diagnosing system errors	Accuracy for fault detection, sensitivity to anomalies	Strategic positioning for fault-prone areas	Fault trend analysis, anomaly detection algorithms	Periodic calibration, low power consumption
Control	Automation, energy management, optimisation	Optimising HVAC systems, adjusting power usage dynamically	Sensors capable of real-time adjustments	Placement for real-time feedback integration	Real-time control feedback loops, optimisation algorithms	Seamless integration with control systems
Monitoring and Reporting	Environmental compliance, data logging, and user interfaces	Logging air quality data, monitoring production trends	Sensors with wide environmental parameter ranges	Distributed placement for comprehensive monitoring	Data aggregation, compliance reports, and visualisation	Reliable operation under varying conditions
Safety and Security	Surveillance, access control, safety assurance	Securing restricted areas, detecting unauthorised entry	High-security sensors, tamper-proof capabilities	Coverage for all critical and high-risk areas	Event-driven processing, security breach detection	Resilient operation in high-security environments
System and Process Optimisation	Efficiency improvement, operational reliability	Streamlining production lines, reducing energy wastage	Energy-efficient, long-lasting, and multi-functional sensors	Optimised placement for process and energy efficiency	Efficiency metrics, operational insights, optimisation models	Sustained performance under demanding conditions

Table 3-1 Sensor Application Domains Comparison

3.3.1 Problem Formulation and Decision Variables

Let **x** = $[x_1, x_2, \ldots, x]^T$ be the binary decision vector, where n is the total number of sensors available. Each element $x_i \in \{0, 1\}$ indicates whether sensor i is selected ($x_i = 1$) or not ($x_i = 0$). A key constraint requires that exactly two sensors must be chosen, which is expressed mathematically as:

$$\Sigma \text{ (from } i = 1 \text{ to } n) \; x_i = 2 \quad \textbf{(3-5)}$$

This binary formulation allows a direct representation of sensor configurations and enables the use of evolutionary algorithms for optimisation.

3.3.2 Objective Functions

The framework defines four objectives that together capture the multifaceted performance requirements. These objectives are combined into a vector-valued function,

$$F(x) = [f_1(x), f_2(x), f_3(x), f_4(x)] \quad \textbf{(3-6)}$$

Where:

Performance (f_1):

This objective is a weighted combination of normalised NDCI and an additional performance metric (for example, accuracy in the OEM scenario or a blend of normalised NDCI and coverage for Airlines and MRO). It is expressed as:

$$f_1(x) = -[\, w_1 \cdot (\Sigma\, x_i \cdot NDCI_i / \max(NDCI)) + w_2 \cdot (\Sigma\, x_i \cdot p_i / \max(p))] \quad \textbf{(3-7)}$$

Where p_i represents the secondary performance metric (such as accuracy), w_1 and w_2 are weights with $w_1 + w_2 = 1$, and the negative sign indicates that the objective is minimised.

Cost (f_2):

The cost objective is defined as the total cost of the selected sensors:

$$f_2(x) = \Sigma\, (x_i \cdot c_i) \quad \textbf{(3-8)}$$

where c_i is the cost of sensor i.

Reliability/Efficiency (f_3):

This objective promotes robust sensor performance by being defined as the negative of a normalised reliability (or efficiency) measure:

$$f_3(x) = -(\Sigma x_i \cdot r_i / \max(r)) \quad \textbf{(3-9)}$$

where r_i denotes the reliability (e.g., mean time between failures) or efficiency of sensor i.

Secondary Criterion (f_4):

For OEM applications, this typically represents the negative of normalised compatibility; for Airlines and MROs, a benefit-to-cost ratio is adopted. For example, in the OEM scenario:

$$f_4(x) = -(\Sigma x_i \cdot d_i / \max(d)) \quad \textbf{(3-10)}$$

where d_i represents the compatibility of sensor i. In other scenarios, f_4 may be defined as:

$$f_4(x) = -(\Sigma x_i \cdot q_i) / (\Sigma x_i \cdot c_i) \quad \textbf{(3-11)}$$

with q_i being a composite benefit measure (such as the sum of normalised NDCI, coverage, and efficiency).

These four objectives collectively express the trade-offs among diagnostic performance, cost, reliability (or efficiency), and secondary considerations.

3.3.3 Optimisation Setup

The optimisation problem is subject to several constraints that ensure system viability:

Binary Selection Constraint:

Each sensor decision variable must be binary, i.e.,
$x_i \in \{0, 1\}$ for all i = 1, …, n, with
Σ (from i = 1 to n), $x_i = 2$.

Budget Constraint:

The total cost of the selected sensors must not exceed a predefined budget, B:

$$\Sigma\ (x_i \cdot c_i) \leq B \tag{3-12}$$

Performance Thresholds:

Minimum acceptable levels are imposed on key metrics:

- The sum of NDCI values must be at least T_{ndci}:

$$\Sigma\ (x_i \cdot NDCI_i) \geq T_{ndci} \tag{3-13}$$

- The average accuracy (or equivalent performance metric) must be at least A_{min}:

$$(\Sigma\ (x_i \cdot p_i)) / (\Sigma\ x_i) \geq A_{min} \tag{3-14}$$

- The total reliability (or efficiency) must be at least R_{min}:

$$\Sigma\ (x_i \cdot r_i) \geq R_{min} \tag{3-15}$$

A MOGA is employed to solve this optimisation problem. MATLAB's "gamultiobj" function is used with a custom creation function that ensures every candidate solution in the initial population consists of exactly two selected sensors. Custom mutation and crossover functions are designed to maintain feasibility (i.e., preserving the constraint $\Sigma\ x_i = 2$) throughout the evolutionary process. The algorithm iteratively evolves the population until convergence is reached, resulting in a Pareto front of non-dominated solutions that represent the best trade-offs among the objectives.

The resulting Pareto set provides a diverse array of sensor-pair configurations. For visualisation, pairwise scatter plots are generated for every combination of two objectives. Each plotted point represents a sensor pair and is annotated with the corresponding sensor names. Feasible solutions are distinctly marked with blue circles, infeasible solutions with red crosses, and the selected optimal pair is highlighted with a green star. This comprehensive analysis supports informed decision-making in complex sensor network implementations and demonstrates the benefits of integrating NDCI into the MOSOF framework.

3.4 Experimental Results and Analysis

3.4.1 Simulation Setup and Parameter Specifications

The proposed NDCI-MOSOF integration was evaluated using simulation studies for three representative scenarios: OEM, Airlines, and MRO. For each case, the sensor pool comprised nine sensors with ECS simulation names and NDCI values. In the OEM scenario, fixed values were used for key parameters (e.g., cost, accuracy, MTBF, and compatibility). In contrast, for the Airlines and MRO cases, the additional attributes (cost, coverage, reliability or efficiency) were generated randomly within realistic bounds to induce nontrivial trade-offs. For instance, in the Airlines scenario, cost was varied between 240 and 280 units, coverage ranged from 0.90 to 0.98, and reliability was randomised between 5100 and 5700. Similar ranges were used for the MRO case, with efficiency values sampled uniformly between 0.80 and 0.95.

MOGA was configured with a population size of 200, a maximum of 500 generations, and a Pareto fraction of 35%. A custom creation function ensured that every individual in the initial population corresponded to exactly two sensors. Other GA parameters, such as mutation (using a feasible-adaptive strategy) and intermediate crossover, were tuned to maintain diversity and feasibility of solutions. Constraints were imposed on total cost, minimum sum of NDCI values, average performance (e.g., accuracy or coverage), and aggregated reliability/efficiency. These constraints were slightly relaxed to ensure a rich solution space, with several sensor pairs meeting all criteria.

3.4.2 Pareto Front Analysis and Sensor Pair Configurations

The optimisation produced a diverse Pareto front capturing the trade-offs among the four objectives: performance (a combination of normalised NDCI and a secondary metric), cost, reliability/efficiency, and a secondary benefit-to-cost or compatibility measure. Detailed pairwise scatter plots (six subplots for every combination of two objectives) reveal that sensor pair solutions span a wide range of trade-offs.

The simulation data was generated with the Boeing 737-800 ECS through the SESAC platform. The ECS dataset is initially organised as a multi-scenario Excel file (e.g., SESAC1.xlsx), where each sheet represents a distinct ECS fault scenario. Each scenario sheet includes a healthy baseline row together with multiple faulted rows, allowing the analysis to compare normal and degraded operation directly. The stored values include sensor readings, baseline values, and a degradation indicator, which is used to represent fault severity in normalised form. As illustrated in Figure 3-8, this raw dataset serves as the Input Data for the NDCI-MOSOF pipeline, which then outlines the complete workflow for ECS sensor optimisation, progressing from preprocessing and fault labelling through to diagnostic scoring, multi-objective search, and the extraction of Pareto-optimal outputs.

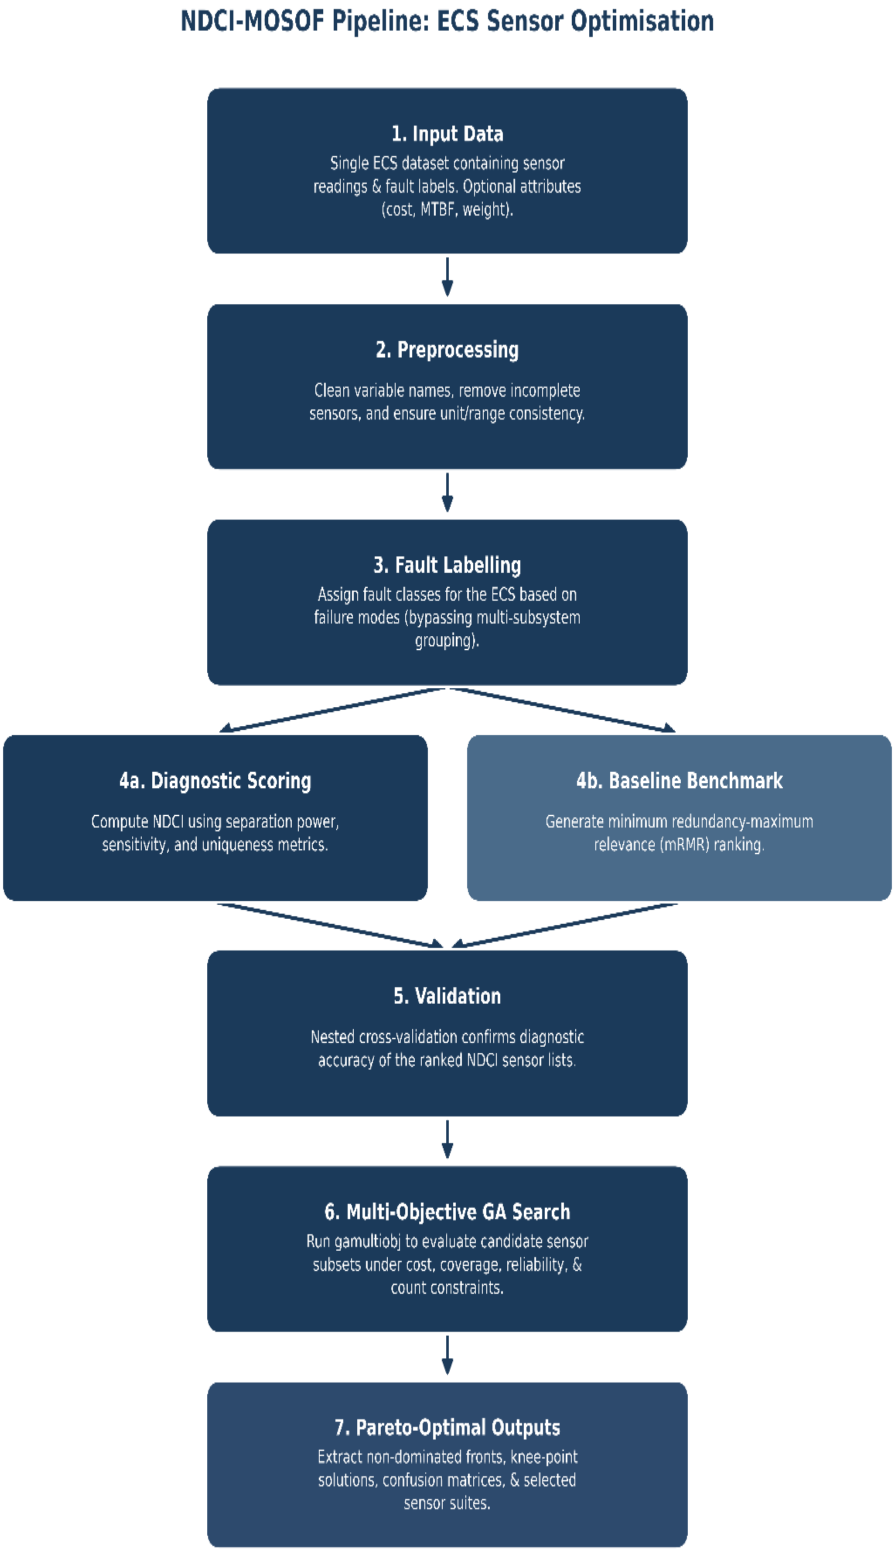

Figure 3-8 NDCI MOSOF pipeline in ECS case

The ECS analysis is severity-aware rather than binary-only. Four ECS fault modes are explicitly analysed: PHX fault, SHX fault, ACM fault, and ram air (RAM) fault. The NDCI formulation uses the healthy baseline to compute SP, scales deviations by the degradation indicator to compute S, and compares sensor responses across the sensor set to compute U. Scenario-level NDCI values are then aggregated across the four ECS cases to rank sensors by their diagnostic contribution.

The data used in this chapter therefore consist of scenario-based ECS sensor responses under progressive degradation, rather than a simple healthy/fault label set. This should be stated explicitly because the diagnostic objective in Chapter 3 is not only fault detection, but also differentiation between ECS fault mechanisms and sensitivity to severity progression. The associated MATLAB workflow referenced in the published paper loads these scenario sheets, computes SP, S, and U, and compares NDCI-based rankings against an mRMR baseline.

Simulation results of each scenario are presented in Figures 3-9, 3-10 and 3-11. MATLAB code of the simulations is available on GitHub [30], facilitating reproducibility and further exploration of the data. Each point in these plots represents a sensor pair, annotated with sensor names. In the OEM scenario, for example, sensor pairs such as "ThiSHX (Temperature, hot, inlet Secondary Heat Exchanger), TiT (Temperature, inlet Turbine)" and "ThiRHX (Temperature, hot, inlet Reheater), P_O (Pack Outlet)" emerged as promising candidates. Feasible solutions (those that met all constraints) are clearly distinguishable from infeasible ones, with blue markers denoting feasible sensor pairs and red crosses indicating violations. The optimal pair (selected as the first feasible solution with the best trade-off among objectives) is highlighted with a green star. Notably, the incorporation of NDCI as a central performance metric enabled greater discrimination between sensor pairs, revealing non-obvious combinations that offer balanced performance and cost advantages. A comparative analysis of the Pareto topologies across the figures reveals the distinct mathematical signatures of the differing stakeholder constraints.

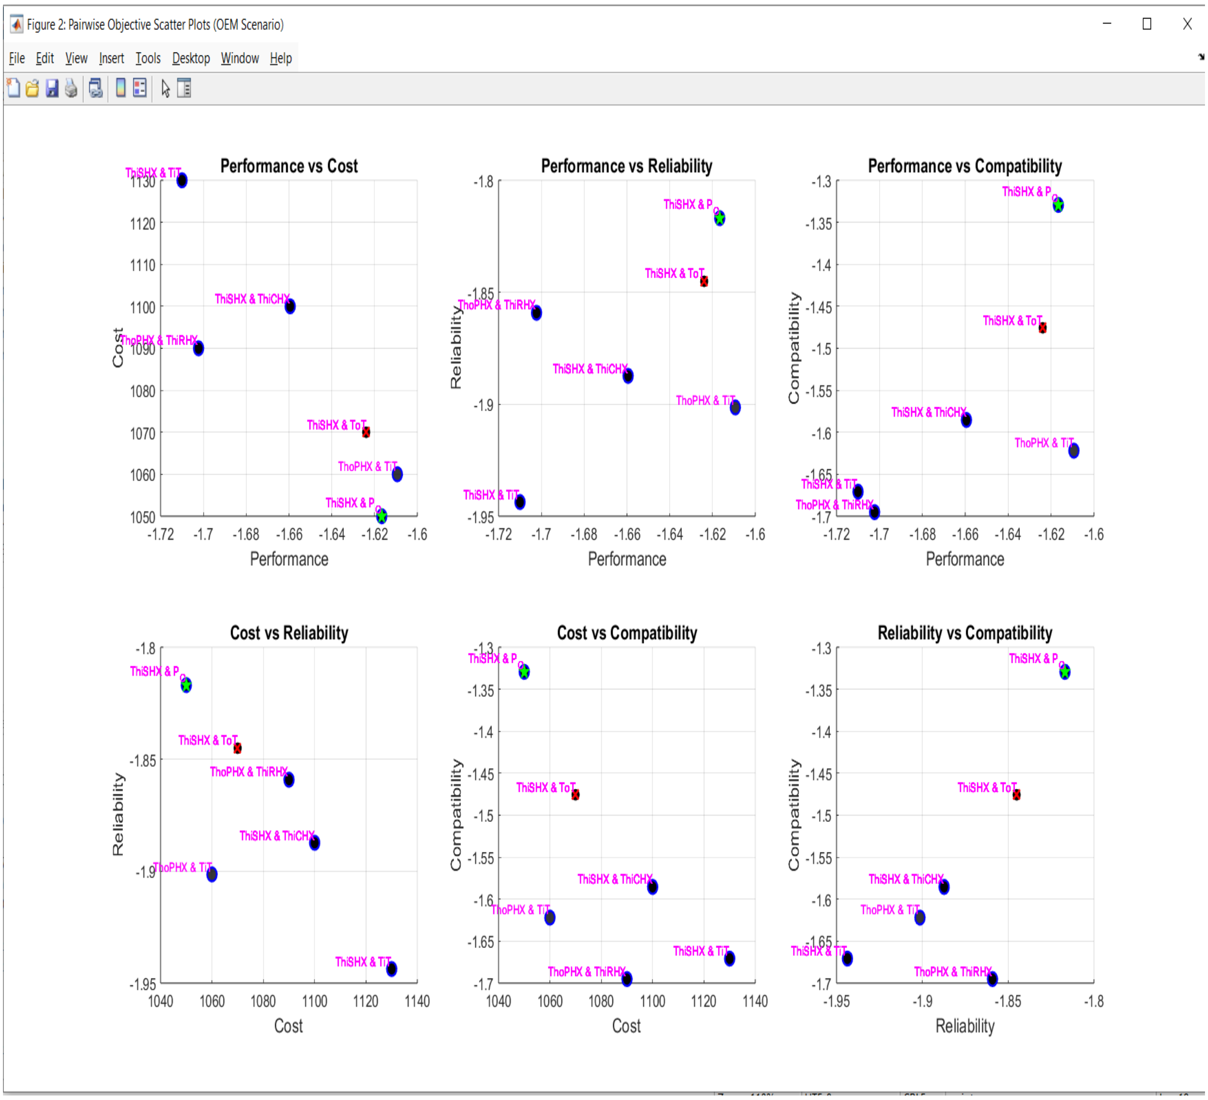

Figure 3-9 OEM case

In the OEM Scenario (Figure 3-9), the objective space is governed by fixed deterministic bounds for certification and unit cost. The resulting scatter plots demonstrate a highly clustered, rigid Pareto front. Because the OEM prioritises strict architectural compatibility and low acquisition cost, the algorithm actively penalises high-performing but novel (or expensive) sensor pairs, forcing the feasible 'green star' solutions into a narrow operational envelope where legacy sensors like ThiSHX and TiT dominate due to their proven certifiability.

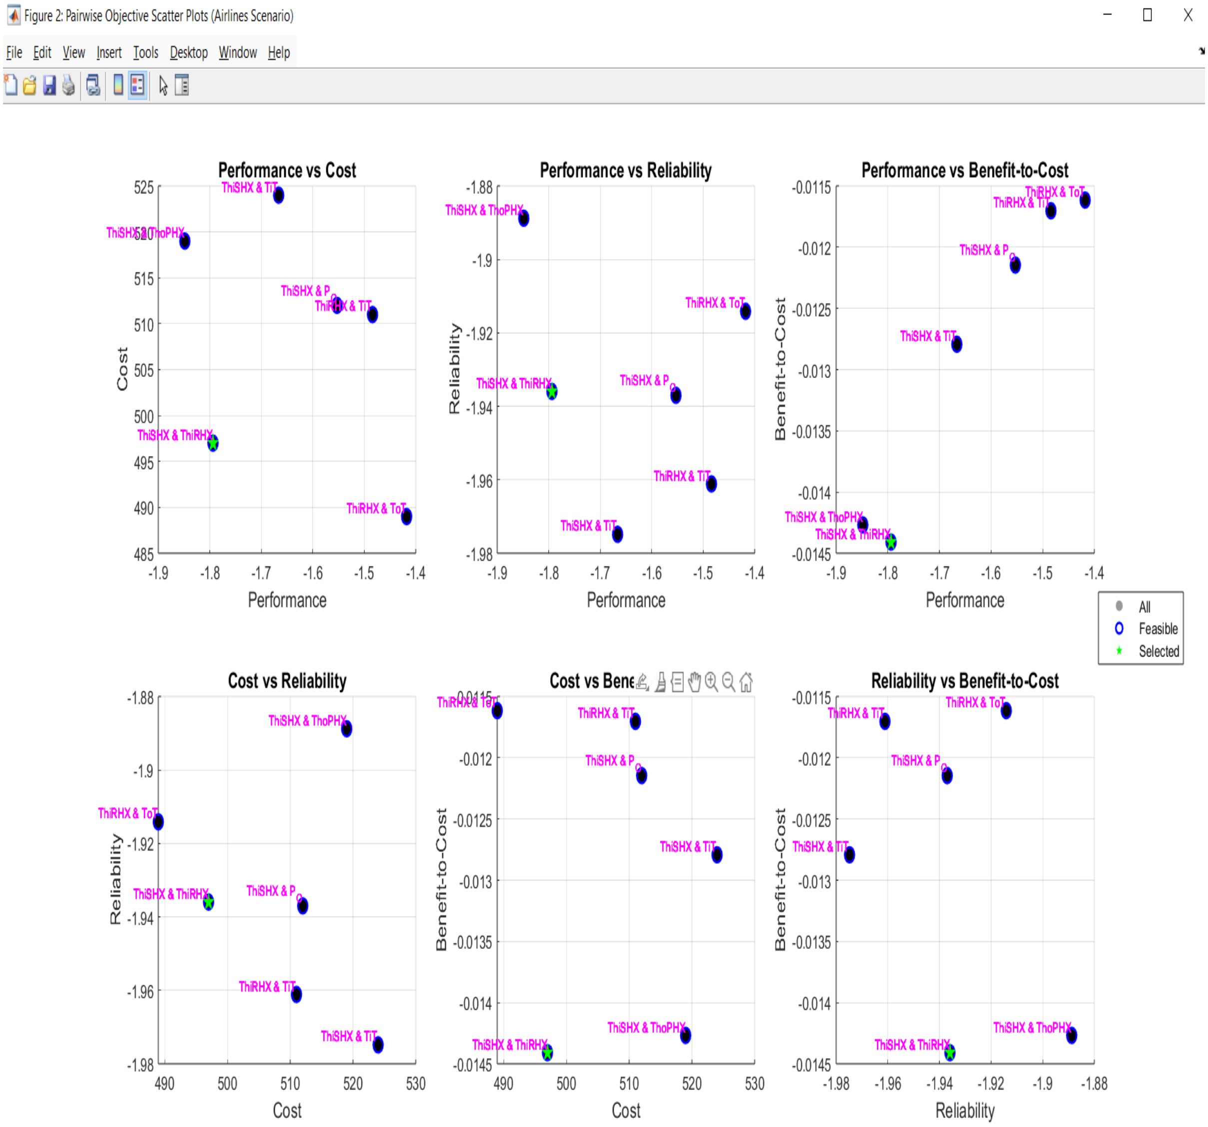

Figure 3-10 Airlines case

Conversely, the Airlines Scenario (Figure 3-10) introduces stochastic variability into the lifecycle cost and reliability bounds, reflecting the uncertainties of a 20-year operational horizon. Consequently, the Pareto front in Figure 3-10 is significantly more dispersed. The 'Performance vs. Benefit-to-Cost' and 'Cost vs. Reliability' subplots illustrate a wider trade-space where the algorithm accepts higher initial unit costs if the sensor pair (e.g., ThiRHX and ToT) guarantees superior long-term reliability and fleet availability.

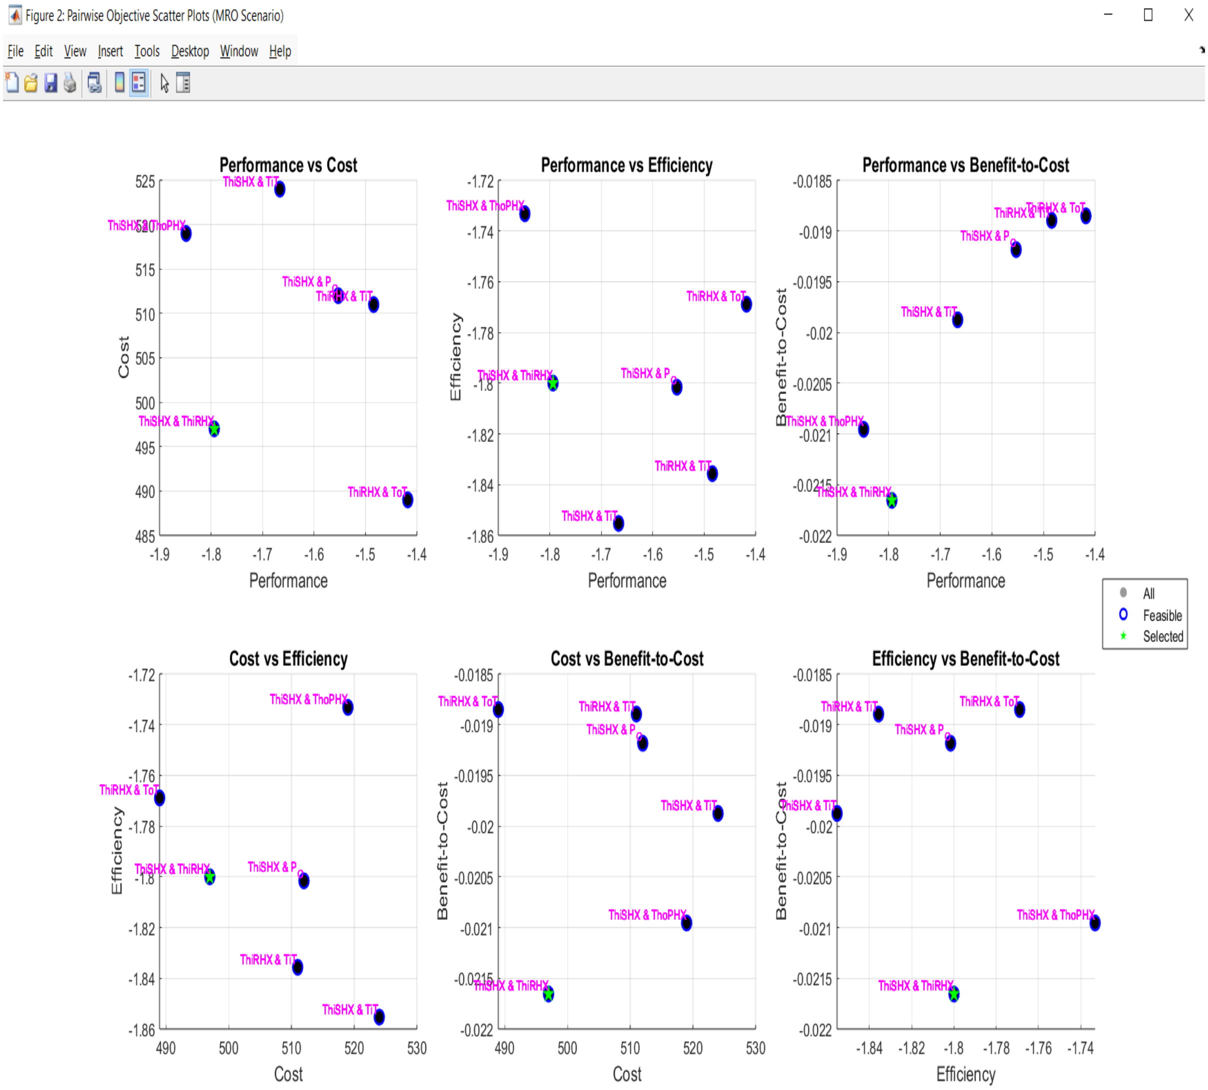

Figure 3-11 MRO Case

Finally, the MRO Scenario (Figure 3-11) shifts the optimisation bias toward rapid fault isolation efficiency. The scatter plots here uniquely highlight the 'Performance vs. Efficiency' trade-off. In this domain, the algorithm heavily favours sensor combinations that yield high NDCI Uniqueness scores, as orthogonal fault signatures directly reduce maintenance troubleshooting time. The dispersion of feasible solutions in Figure 3-11 confirms that when turnaround time is the primary economic driver, MROs are mathematically justified in deploying high-efficiency sensor architectures that might be deemed overly expensive in a strict OEM unit-cost paradigm. Ultimately, these three figures collectively validate the core hypothesis of the MOSOF framework: that an 'optimal' sensor suite is inherently relative, and its architectural topology shifts dynamically based on the economic and temporal horizon of the evaluating stakeholder.

3.5 Discussion

The experimental results demonstrate that integrating the NDCI into the MOSOF framework enhances sensor selection in diagnostic applications for complex systems by consolidating multiple aspects and different users' objectives into a single framework for sensor design implementations. In contrast to the conventional MOSOF, the NDCI-MOSOF approach introduces a performance metric that captures subtle yet critical differences among sensor outputs. This improved discrimination is evident in the broader Pareto spread and more diversified sensor pair configurations obtained in the simulations.

One major implication is that the NDCI-centric framework facilitates a more balanced trade-off between performance and cost. By incorporating NDCI as a core metric in the performance objective, the framework successfully identifies sensor pairs that are not only cost-effective but also exhibit superior sensitivity and fault detection capabilities. This is particularly relevant in applications such as control systems, process monitoring, and security, where high-fidelity sensor performance is essential.

Furthermore, the adaptability of the proposed method to different application scenarios (OEM, Airlines, MRO) underscores its versatility. The random scattering of secondary attributes in the Airlines and MRO cases further validates the framework's robustness against non-ideal correlations among sensor attributes. The ability to adjust constraints and objective weightings makes the framework suitable for a wide range of sensor implementations, ensuring that system-specific requirements are met.

However, the study also highlights potential limitations, such as the reliance on simulation data. Future research should focus on validating the NDCI-MOSOF approach with empirical data. Additionally, extending the framework to incorporate dynamic sensor behaviour and real-time adaptation remains a promising research direction.

3.6 Conclusion

This study presents a novel integration of the NDCI into the MOSOF to address the nuanced performance requirements of complex sensor systems. By embedding NDCI as a central performance metric and augmenting the framework with additional objectives (such as cost, reliability/efficiency, and secondary factors like compatibility or benefit-to-cost ratio), the proposed NDCI-MOSOF approach provides a more precise and application-specific sensor selection process.

Simulation studies across OEM, Airlines, and MRO scenarios demonstrate that the enhanced framework can generate a diversified Pareto front of feasible sensor pair configurations. These results indicate improved discrimination among sensor pairs, leading to a better balance between performance and economic considerations. The approach also shows strong potential for adaptation to other complex system diagnostic applications.

In conclusion, integrating NDCI into MOSOF represents a significant advancement in sensor optimisation methodologies. The enhanced framework not only refines sensor selection by focusing on critical performance attributes but also offers a flexible tool for tailoring sensor configurations to diverse and dynamic application environments. Future work will focus on real-world validation and dynamic adaptation to further extend the applicability of the proposed method.

3.7 References

1. G. Vachtsevanos, F. Lewis, M. Roemer, A. Hess, and B. Wu, Intelligent Fault Diagnosis and Prognosis for Engineering Systems. 2007. doi:10.1002/9780470117842
2. Y. Lei, N. Li, L. Guo, N. Li, T. Yan, and J. Lin, "Machinery health prognostics: A systematic review from data acquisition to RUL prediction," 2018. doi: 10.1016/j.ymssp.2017.11.016.
3. H. Peng, F. Long, and C. Ding, "Feature selection based on mutual information: Criteria of Max-Dependency, Max-Relevance, and Min-Redundancy," *IEEE Trans Pattern Anal Mach Intell*, vol. 27, no. 8, 2005, doi: 10.1109/TPAMI.2005.159.
4. M. Aljawarneh, R. Hamdaoui, A. Zouinkhi, B. Boussaid, and M. N. Abdelkrim, "Energy efficiency approaches in Wireless Sensor Networks," in *2022 IEEE 21st International Conference on Sciences and Techniques of Automatic Control and Computer Engineering, STA 2022 - Proceedings*, 2022. doi: 10.1109/STA56120.2022.10019187.
5. Jo, S. Lee, and S. Oh, "Improved measures of redundancy and relevance for mRMR feature selection," *Computers*, vol. 8, no. 2, 2019, doi: 10.3390/computers8020042.
6. L. M. Santi, T. S. Sowers, and R. B. Aguilar, "Optimal sensor selection for health monitoring systems," in *41st AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit*, 2005. doi: 10.2514/6.2005-4485.
7. S.-M. Yang, J. Qiu, G.-J. Liu, and P. Yang, "Sensor Selection and Optimization for Aerospace System Health Management under Uncertainty Testing."
8. M. Esperon-Miguez, I. K. Jennions, I. C. Escobar, and N. Hanov, "Simulating faults in a boeing 737-200 environmental control system using a thermodynamic model," *Int J Progn Health Manag*, vol. 10, no. 2, 2019, doi: 10.36001/ijphm.2019.v10i2.2731.
9. L. Jia, C. M. Ezhilarasu, and I. K. Jennions, "Cross-Condition Fault Diagnosis of an Aircraft Environmental Control System (ECS) by Transfer Learning," *Applied Sciences (Switzerland)*, vol. 13, no. 24, 2023, doi: 10.3390/app132413120.
10. Z. Xu, Y. Guo, and J. Homer Saleh, "Multi-objective optimization for sensor placement: An integrated combinatorial approach with reduced order model and Gaussian process," *Measurement (Lond)*, vol. 187, 2022, doi: 10.1016/j.measurement.2021.110370.
11. Lee *et al.*, "Multi-objective instance weighting-based deep transfer learning network for intelligent fault diagnosis," *Applied Sciences (Switzerland)*, vol. 11, no. 5, 2021, doi: 10.3390/app11052370.
12. Kabashkin, V. Perekrestov, T. Tyncherov, L. Shoshin, and V. Susanin, "Framework for Integration of Health Monitoring Systems in Life Cycle Management for Aviation Sustainability and Cost Efficiency," *Sustainability 2024, Vol. 16, Page 6154*, vol. 16, no. 14, p. 6154, Jul. 2024, doi: 10.3390/SU16146154.
13. B. Suslu, F. Ali, and I. K. Jennions, "Understanding the Role of Sensor Optimisation in Complex Systems," 2023. doi: 10.3390/s23187819.
14. J. Xu, Y. Wang, and L. Xu, "PHM-Oriented Sensor Optimization

Selection Based on Multiobjective Model for Aircraft Engines," *IEEE Sens J*, vol. 15, no. 9, pp. 4836–4844, Sep. 2015, doi: 10.1109/JSEN.2015.2430361.

15. C. A. Farschman, K. P. Viswanath, and B. E. Ydstie, "Process systems and inventory control," *AIChE Journal*, vol. 44, no. 8, pp. 1841–1857, Aug. 1998, doi: 10.1002/AIC.690440814.
16. E. Clark, S. L. Brunton, and J. N. Kutz, "Multi-Fidelity Sensor Selection: Greedy Algorithms to Place Cheap and Expensive Sensors with Cost Constraints," *IEEE Sens J*, vol. 21, no. 1, pp. 600–611, Jan. 2021, doi: 10.1109/JSEN.2020.3013094.
17. M. Civera, M. L. Pecorelli, R. Ceravolo, C. Surace, and L. Zanotti Fragonara, "A multi-objective genetic algorithm strategy for robust optimal sensor placement," *Computer-Aided Civil and Infrastructure Engineering*, vol. 36, no. 9, pp. 1185–1202, Sep. 2021, doi: 10.1111/mice.12646.
18. J. Hare, S. Gupta, N. Najjar, P. D'Orlando, and R. Walthall, "System-Level Fault Diagnosis with Application to the Environmental Control System of an Aircraft," in *SAE Technical Papers*, 2015. doi: 10.4271/2015-01-2583.
19. F. M. de S. Mello and G. F. Gomes, "Sensor placement optimization for composite aircraft structures: A multi-objective Kriging-based approach," *Compos Struct*, vol. 353, p. 118723, Jan. 2025, doi: 10.1016/J.COMPSTRUCT.2024.118723.
20. "New Aircraft Cabin Environment Sensor Certified for Boeing 737 - Avionics International." Accessed: Mar. 09, 2025. [Online]. Available: https://www.aviationtoday.com/2021/03/23/new-aircraft-cabin-environment-sensor-certified-boeing-737/
21. C. M. Ezhilarasu, Z. Skaf, and I. K. Jennions, "A Generalised Methodology for the Diagnosis of Aircraft Systems," *IEEE Access*, vol. 9, pp. 11437–11454, 2021, doi: 10.1109/ACCESS.2021.3050877.
22. L. Mao, B. Davies, and L. Jackson, "Application of the Sensor Selection Approach in Polymer Electrolyte Membrane Fuel Cell Prognostics and Health Management," *Energies 2017, Vol. 10, Page 1511*, vol. 10, no. 10, p. 1511, Sep. 2017, doi: 10.3390/EN10101511.
23. D. Lee, I. Lee, Y. Kim, S. C. Joo, and J. H. Choi, "Sensor set optimization by functional model and Bayesian network for fault diagnosis of turbine generator lubrication system," *Eng Appl Artif Intell*, vol. 138, p. 109416, Dec. 2024, doi: 10.1016/J.ENGAPPAI.2024.109416.
24. T. Planès, S. Delbecq, V. Pommier-Budinger, and E. Bénard, "Modeling and Design Optimization of an Electric Environmental Control System for Commercial Passenger Aircraft," *Aerospace*, vol. 10, no. 3, 2023, doi: 10.3390/aerospace10030260.
25. Q. Liu, L. Zhuang, J. Wen, B. Dong, and Z. Liu, "Thermodynamic Optimization of Aircraft Environmental Control System Using Modified Genetic Algorithm," *Processes*, vol. 10, no. 4, 2022, doi: 10.3390/pr10040721.
26. J. M., S. J. E. House, "SYSTEMS AND METHODS FOR FAULT DETECTION OF AIR HANDLING UNITS," US 8.239,168 B2, 2012

27. Q. "KUMAGAI, G. " D. B. "SHAO, “PATTERN CLASSIFICATION SYSTEM WITH SMART DATA COLLECTION FOR ENVIRONMENTAL CONTROL SYSTEM FAULT ISOLATION,” EP 3 719 602 B1, 2020
28. Jennions and F. Ali, “Assessment of heat exchanger degradation in a Boeing 737-800 environmental control system,” *J Therm Sci Eng Appl*, vol. 13, no. 6, 2021, doi: 10.1115/1.4050324.
29. B.Suslu-“NDCI_Calculations,”https://github.com/ssl8/ECS-sensor-optimisation/blob/4174e20aff3e1c189d72aa9c75fbd104a5c339c7/multiFaultComparison.m.
30. B.Suslu-“MOSOF_ECS_Case_Outlines,”https://github.com/ssl8/ECS-sensor-optimisation/blob/0ce6ca4df91cb1fc3f2c98b2483d4b9b6caed0d9/NDCI_MOSOF_ECS.m.

4 MOSOF with NDCI: A Cross-Subsystem Evaluation of an Aircraft for an Airline Case Scenario

Chapter 4 (based on the third article) extends the NDCI–MOSOF framework beyond a single subsystem, applying it across four major aircraft subsystems – the Engine, Fuel System (FS), Electrical Power System (EPS), and Environmental Control System (ECS). This chapter demonstrates the framework's scalability and system-of-systems perspective by optimising sensor networks in a coordinated, platform-level context. Building on the methodology validated in Chapter 3, each subsystem is represented by a high-fidelity simulation model that encompasses multiple fault modes and severity levels. The data from these models is integrated to capture cross-subsystem fault propagation, ensuring that sensor selection accounts for interdependencies between systems. The Normalised Diagnostic Contribution Index (NDCI) is employed to rank sensor signals by diagnostic importance in each scenario. A multi-objective genetic algorithm (MOGA) then searches for Pareto-optimal sensor suites that balance diagnostic performance with practical constraints (such as cost, weight, and reliability) and diverse stakeholder priorities (OEM, airline, MRO requirements).

Key findings: The results show that the enhanced NDCI–MOSOF framework successfully optimises sensor configurations for all four subsystems simultaneously, achieving high fault detection rates and accurate isolation of failures with a relatively minimal sensor set. By considering all subsystems together, the approach captures inter-subsystem fault effects that would be missed if each system were optimised in isolation. For instance, the optimisation identified sensor sets that not only detect faults within individual systems but also quickly recognise faults cascading between the electrical and fuel systems. The NDCI-driven solutions consistently outperformed a baseline mRMR-based selection, delivering comparable or better diagnostic coverage with fewer sensors, all while satisfying the imposed cost and reliability constraints. Overall, Chapter 4 confirms the generality and versatility of the proposed framework: it is not a one-off solution tailored to a single subsystem, but a broadly applicable strategy for whole-aircraft health monitoring. This multi-subsystem demonstration marks a significant step toward comprehensive, platform-wide IVHM sensor optimisation, providing quantitative evidence that a holistic optimisation approach can improve fault diagnosis across an integrated aircraft system.

4.1 Introduction

The advancement of Integrated Vehicle Health Management (IVHM) is crucial to enhancing the safety, reliability, and operational efficiency of modern, complex assets, such as aircraft [1]. Modern aircraft have evolved from collections of discrete components into highly integrated systems-of-systems, where mechanical, thermal, electrical, and computational subsystems are deeply interconnected. A fundamental challenge in designing effective IVHM systems lies in the strategic selection of sensors [2]. The conventional approach of instrumenting systems primarily for control is often insufficient for comprehensive diagnostics, as the subtle signatures of incipient faults can be easily missed [3]. Faults in highly integrated systems rarely manifest in isolation. For instance, a malfunction in an engine's bleed air system can propagate, causing performance shifts in the Environmental Control System (ECS) and load fluctuations in the Electrical Power System (EPS) [4,5]. An optimal sensor network must therefore possess the capability to capture these distributed, cross-subsystem effects, enabling accurate fault isolation [6,7].

The process of selecting this optimal sensor set is a complex task involving dimensionality reduction. While generic feature selection algorithms, such as Minimum Redundancy–Maximum Relevance (mRMR) [8], are widely used, they have inherent limitations in the context of engineering diagnostics [9]. Such methods typically rely on statistical relevance and redundancy criteria, which do not fully capture a sensor's true diagnostic utility. They may fail to differentiate between faults of varying severity and can discard correlated sensors that, despite statistical similarity, measure distinct physical phenomena crucial for robust diagnosis. This recognition led to foundational work establishing the need for domain-aware metrics that evaluate sensors based on their contribution to diagnostic objectives [2].

This paper builds directly upon that foundation. The previous research first conceptualised the limitations of generic methods and proposed a diagnostically focused evaluation framework [1]. Then, the Normalised Diagnostic Contribution Index (NDCI) is introduced, demonstrating its initial application within a Multi-Objective Sensor Optimisation Framework (MOSOF) on a single aircraft subsystem [2]. The present study provides a comprehensive, cross-subsystem validation and refinement of this integrated NDCI-MOSOF methodology.

Sensor selection is intrinsically a multi-objective problem, requiring a concurrent optimisation of often-conflicting goals. Stakeholders, including operators, manufacturers, and maintenance providers, must navigate the trade-space between maximising diagnostic coverage and minimising factors such as instrumentation cost, weight, and reliability penalties [10,11]. Single-objective optimisation cannot adequately address this complex decision-making landscape. Therefore, the MOSOF approach leverages a multi-objective genetic algorithm to explore this trade-space, identifying a set of non-dominated, or Pareto-optimal, solutions [12,13]. This methodology provides decision-makers with a portfolio of optimised sensor suites, each representing a different balance of priorities. This aligns with modern aerospace development guidelines, like ARP4754B [14], which advocate for systematic, model-based processes to substantiate safety and operational requirements [15].

Navigating this complex trade-space is a central challenge in IVHM design.

The trade-offs involved are significant and have direct operational consequences. The true cost of an aerospace-grade sensor extends far beyond its unit manufacturing price, including substantial overheads for certification, material traceability, regulatory documentation, and lifecycle support infrastructure [16]. It is not uncommon for a sensor with a manufacturing cost of €500 to have a market price of €5,000 once these factors are included [16]. Similarly, weight is a critical consideration in aircraft design. Every kilogram of mass added to an airframe translates into significant lifecycle fuel costs, with industry estimates placing the cost of weight between $1,500 and $2,000 per kilogram over the aircraft's operational life [17]. An optimised sensor suite that saves even a few kilograms can therefore yield millions of dollars in operational savings for airlines, as the cost of weight between $100-$1000/kg is confirmed to be reasonable in other studies [18]. Finally, every component added to a system introduces a new potential point of failure, which can decrease the overall system's Mean Time Between Failures (MTBF). This reliability penalty must be carefully weighed against the diagnostic benefits the sensor provides.

This multi-objective landscape is further complicated by the divergent priorities of the key stakeholders involved in an aircraft's lifecycle. Original Equipment Manufacturers (OEMs) are often focused on minimising manufacturing costs, simplifying the design to streamline certification against standards like ARP4754B, and ensuring compliance with airworthiness directives. Airlines and other operators, conversely, prioritise operational continuity, seeking to maximise fleet availability, reduce lifecycle maintenance costs, and achieve the highest possible benefit-to-cost ratio for any new technology they adopt. Finally, Maintenance, Repair, and Overhaul (MRO) providers value sensor systems that enable rapid and unambiguous fault isolation, as this directly reduces aircraft-on-ground (AOG) time and expedites the repair process. An optimal sensor selection framework must be capable of reconciling these varied and often competing interests.

This physical reality of fault propagation demands a holistic, system-of-systems approach to sensor placement. A sensor suite that has been optimised for a single subsystem in isolation is likely to be globally suboptimal and may even be unsafe. Traditional design processes, which often involve separate component teams submitting lists of desired measurements for their specific subsystems, create a siloed view of the aircraft's health [19]. This approach inevitably creates a "diagnostic blind spot" at the interfaces between subsystems. For instance, an optimisation focused solely on the engine might discard a sensor that, while only moderately useful for diagnosing engine-internal faults, is critically important for observing the downstream effects of a bleed air anomaly on the

ECS. Conversely, an ECS-focused optimisation would have no logical basis for selecting a sensor located on the engine, even if that sensor provides the earliest and clearest indication of an impending bleed air problem. Only a unified optimisation framework that considers all candidate sensors and all potential fault modes across the entire platform simultaneously can "see" into this blind spot. Such a framework can identify and select sensors based on their cross-subsystem diagnostic value. This capability is not merely an improvement but a necessary condition for achieving comprehensive fault coverage in modern, integrated aircraft.

This study makes the following contributions to help overcome this bottleneck problem in the diagnostics domain: (i) It benchmarks the refined NDCI against mRMR across four critical aircraft subsystems using a robust nested cross-validation protocol for reliable performance estimation. (ii) It demonstrates the framework's ability to identify and leverage cross-subsystem synergies for improved diagnostics. (iii) It applies the complete MOSOF to a realistic airline case study, identifying a practical, cost-effective sensor suite from the Pareto-optimal front.

4.2 Materials and Methods

This section formalises how the NDCI–MOSOF pipeline, from a single-subsystem study to a platform-level optimisation across four interacting aircraft subsystems—engine, fuel, electrical power system (EPS), and environmental control system (ECS) — is extended. The goal is to select sensor suites that (i) detect degradation, (ii) isolate fault modes, and (iii) respect stakeholder constraints (OEM, Airline, MRO) on cost, reliability/efficiency, and compatibility. The modelling and optimisation machinery, faithful to the previous work and its ECS validation, is retained, and the data interfaces, metrics, and constraints that allow cross-subsystem synergies and scenario-wise coverage are generalised.

The process started by assembling a platform-level data matrix from the four subsystems. Each record is associated with a *scenario* (healthy or a specific fault mode at a given severity). Pre-processing unifies sampling and units, removes outliers and near-constant channels, and applies z-score normalisation/standardisation. A mapping table links every sensor to its owning subsystem and to the set of fault modes it can potentially inform. This yields the *scenario sheets* used throughout the pipeline.

Feature ranking then proceeds along two complementary streams. The first is the NDCI stream, where three ingredients: (1) Separability (*SP*): how well a sensor separates the scenario from competing classes; (2) Sensitivity: monotonic change with degradation severity; and (3) Uniqueness: non-redundancy with respect to the current candidate set are computed per scenario. These are normalised and combined into an NDCI matrix, which is aggregated across scenarios to obtain an overall ranking. In parallel, a baseline mRMR stream produces a global relevance–redundancy ranking using the same labels. The dual stream provides a robustness check and an independent baseline for comparison.

Because a single global ranking can under-serve rare but safety-critical cases, next, the coverage-based minimal subsets per scenario are derived. For each scenario, the smallest subset achieving a pre-specified fraction (e.g., ≥95%) of the best achievable scenario-NDCI is retained. These subsets are computed separately for the NDCI and mRMR streams, and their union is taken across scenarios, a simple technique that preserves coverage without increasing the set size.

The resulting candidate pools serve as the seed for the MOSOF, which is solved using a multi-objective genetic algorithm (MOGA). The decision variable encodes exactly k sensors per policy. Objectives include performance (NDCI and/or classifier coverage/accuracy), cost (Σ cost), reliability/efficiency (e.g., Σ MTBF−1^−1 or Σ efficiency loss), and a stakeholder criterion (OEM compatibility or Airline/MRO benefit-to-cost). The output of the MOGA process is not a single "optimal" solution but rather a set of non-dominated, or Pareto-optimal, solutions. A sensor suite is defined as Pareto-optimal if no single objective (such as diagnostic performance) can be improved without simultaneously degrading at least one other objective (such as cost or reliability) [20]. This collection of non-dominated solutions forms the

Pareto front, which represents the boundary of best-achievable designs. This Pareto front serves as a powerful decision-support tool. It provides stakeholders with a transparent portfolio of optimised sensor suites, each representing a different, quantitatively evaluated trade-off between performance, cost, reliability, and other key objectives. Hard constraints enforce a budget cap, a minimum summed NDCI, a minimum average diagnostic performance, and a minimum reliability/efficiency. The outcome is a Pareto front from which stakeholder-specific choices can be read off.

Finally, the selected suites, using repeated/nested cross-validation with lightweight classifiers (bagged decision tree), are validated and generalised. The detection (balanced accuracy) and isolation (per-fault mode accuracy) are reported, along with confusion matrices, stepwise feature histories, and low-dimensional separability visuals for interpretability.

Figure 4-1 provides a high-level architectural summary of the NDCI–MOSOF pipeline. To make the computational sequence fully reproducible, the operational workflow, simulation settings, and code-module interactions are stated explicitly.

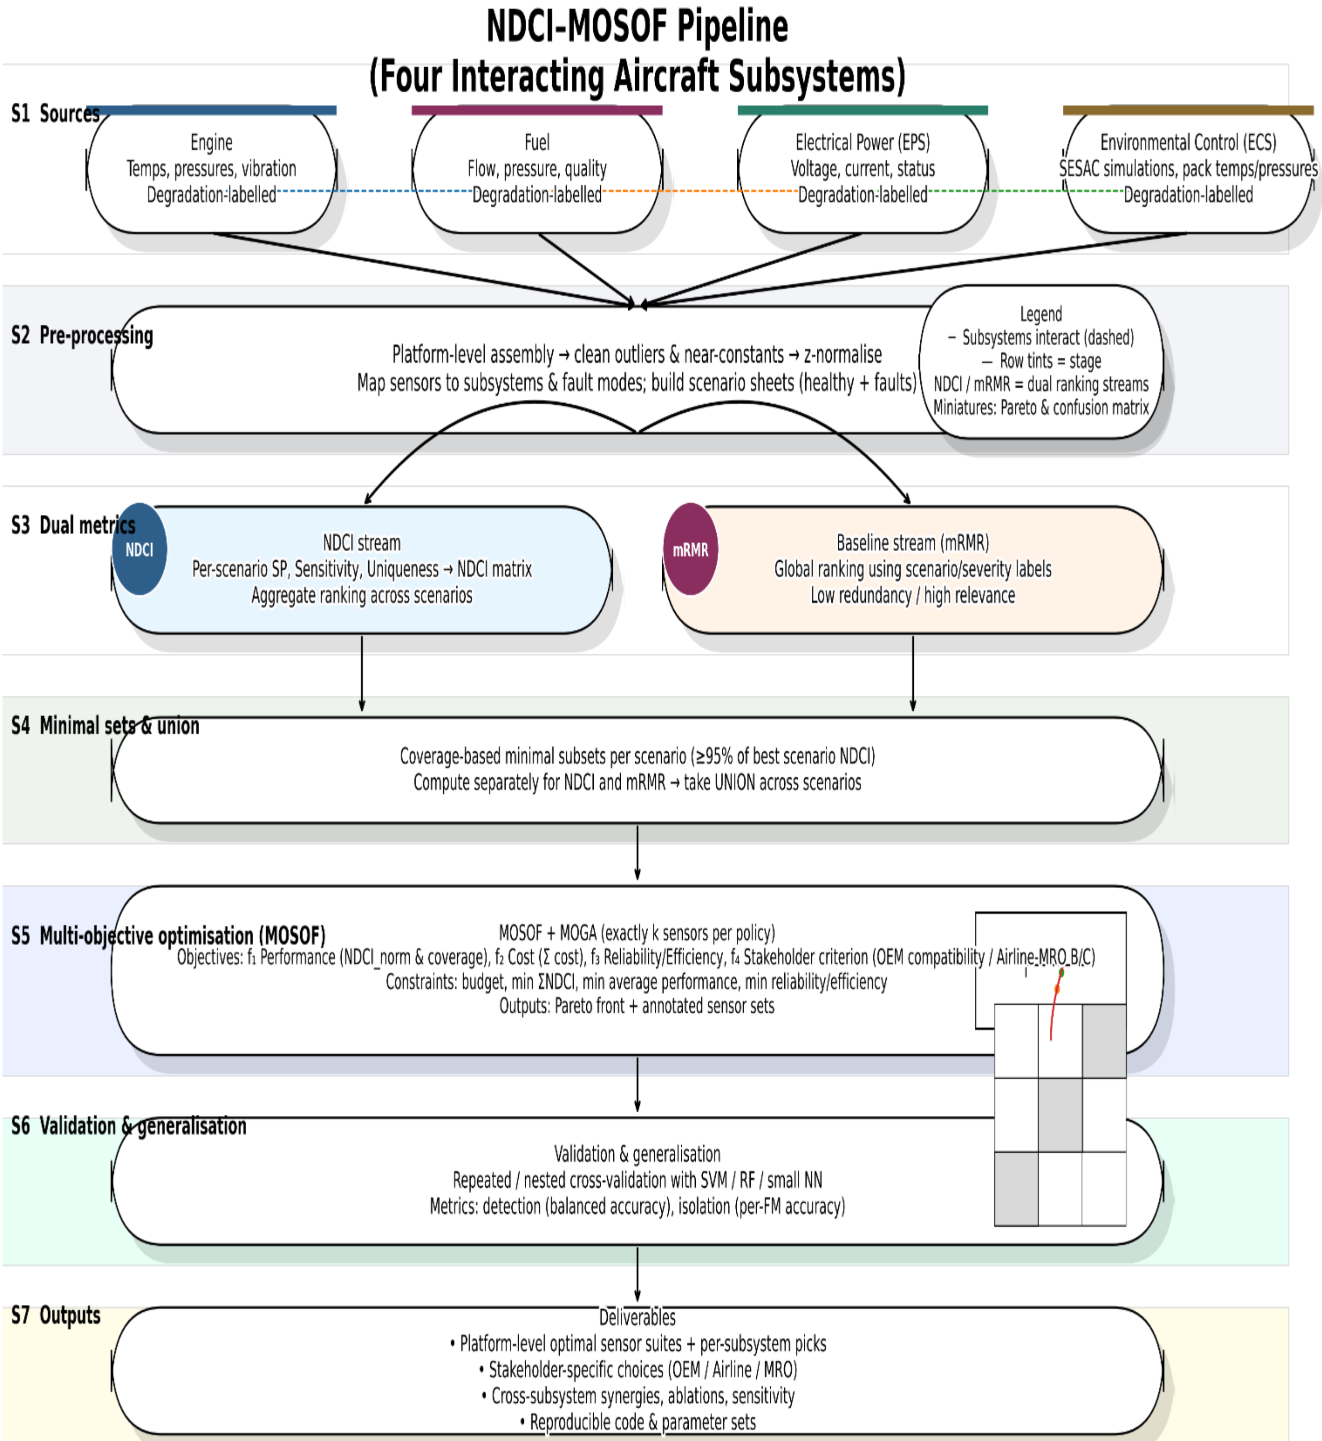

Figure 4-1 NDCI–MOSOF methodology for multi-subsystem aircraft diagnostics

The end-to-end workflow used in this study is summarised in execution order, so that the interaction between the data, ranking modules, optimisation stage, and validation stage is explicit.

- **S1-Sources:** Multi-sensor streams from four interacting subsystems—engine, fuel, electrical power system (EPS), and environmental control system (ECS)—with scenario labels covering healthy and fault-mode/severity conditions; dashed connectors indicate cross-subsystem couplings.
- **S2-Pre-processing: Platform-level** assembly; removal of outliers and near-constant channels; temporal alignment and z-normalisation; mapping of sensors to subsystems and fault modes; construction of scenario sheets.
- **S3-Dual metrics:** (i) *NDCI stream* computes, per scenario, separability (SP), sensitivity to degradation, and class-uniqueness to form an NDCI matrix and aggregate ranking across scenarios; (ii) *mRMR baseline* provides a relevance–redundancy ranking using the same scenario/severity labels.
- **S4-Minimal sets & union:** For each scenario, minimal sensor subsets that retain ≥95% of the best-scenario NDCI are found separately for NDCI and mRMR; unions across scenarios preserve coverage against case variation.
- **S5-MOSOF optimisation:** A multi-objective genetic algorithm (MOGA) searches *exactly k*-sensor policies under constraints (budget cap, minimum ΣNDCI, minimum average detection/isolation performance, minimum reliability/efficiency). Objectives span (f_1) performance (NDCI and/or classifier coverage/accuracy), (f_2) cost, (f_3) reliability/efficiency, and (f_4) a stakeholder criterion (OEM compatibility or Airline/MRO benefit-to-cost). The panel includes a schematic Pareto front for emphasis.
- **S6-Validation & generalisation:** Repeated/nested cross-validation with SVM, Random Forest, and small neural nets; metrics include balanced accuracy for detection and per-fault-mode isolation; artefacts include confusion matrices and low-dimensional separability visuals.
- **S7-Outputs:** Pareto fronts and recommended suites at platform and subsystem level, stakeholder-specific selections (OEM/Airline/MRO),

cross-subsystem synergies, and packaged code/parameters for reproducibility.

A fundamental premise of the vehicle-level diagnostic framework is that modern aircraft are highly integrated systems-of-systems; consequently, localized faults rarely remain confined to their system of origin. The four subsystems evaluated in this study—the Engine, Fuel System, EPS, and ECS—are bound together by strict thermodynamic, fluid, and electro-mechanical physical linkages.

1. Fuel to Engine: Fuel to Engine: The fuel system is responsible for supplying the chemical energy required for propulsion. Anomalies such as fuel pump leakage or Fuel Oil Heat Exchanger (FOHE) clogging reduce the fuel mass flow rate and alter manifold pressures. To maintain commanded thrust, the engine's full authority digital engine control compensates by modulating operation, which inherently alters the engine's internal combustion dynamics and shifts its global thermodynamic signature, manifesting in specific measurable parameters such as Exhaust Gas Temperature (EGT), low-pressure rotor speed (N1), and high-pressure rotor speed (N2).

2. Engine to EPS (Electro-Mechanical Coupling): The engine shaft drives the integrated drive generators; a shaft-speed anomaly therefore propagates as a voltage/frequency deviation in the EPS.

3. Engine to ECS (Pneumatic and Thermal Coupling): Traditional commercial aircraft environmental control systems are entirely reliant on high-pressure, high-temperature bleed air extracted from the engine's compressor stages. Aerodynamic faults within the engine, such as High-Pressure Compressor (HPC) contamination or bleed valve malfunctions, directly perturb the pressure and mass flow of the air delivered to the ECS. This restricts the heat rejection capability of the downstream Primary Heat Exchanger (PHX) and forces the Air Cycle Machine (ACM) to operate off-design, creating a direct pneumatic-to-thermal fault propagation pathway.

4. EPS to ECS (Electro-Mechanical Coupling): The EPS provides the critical power required for actuation and control logic across the platform. A failure in the EPS, such as an intermittent supply to the ECS Temperature Control Valve (TCV), causes the valve to stick. This electrical fault instantly propagates into the thermal domain, allowing unregulated hot bypass air to bypass the cooling cycle, causing severe temperature spikes at the PACK outlet. Conversely, mechanical degradation within the ACM increases torque demands, which translates back into the EPS as abnormal electrical load spikes.

4.2.1 Normalised Diagnostic Contribution Index (NDCI) vs. mRMR

The NDCI is a composite metric specifically engineered to quantify a sensor's diagnostic value, moving beyond generic statistical measures. As established in [2], the NDCI integrates three normalised components that reflect key diagnostic properties:

Separation Power (SP): This component quantifies the magnitude of a sensor's response to a fault relative to its baseline noise. It is calculated as the mean absolute deviation of the sensor signal from its healthy baseline across all fault conditions, normalised by the dynamic range observed during healthy operation.

Severity Sensitivity (S): This component measures how well a sensor's response tracks the progression of a fault. It is computed by normalising the sensor's absolute deviation by a factor representing the fault's severity, rewarding sensors that show a monotonic response as degradation deepens.

Uniqueness (U): This component promotes informational diversity within the sensor suite. To avoid selecting clusters of redundant sensors, it penalises signals that are highly correlated with others. It is computed as one minus the average absolute Pearson correlation between a given sensor and all others under fault conditions.

Each component (SP, S, and U) is scaled to a range and averaged to produce a single, unified NDCI score. A high-NDCI sensor, therefore, exhibits a clear and robust response to faults, tracks fault progression, and provides information that is distinct from other sensors in the network.

FM's descriptions and the complete list of associated subsystems are shown in Table 4-1.

Table 4-1 Fault Mode Descriptions

Fault Mode	Compact physical description
FM1 – AC Motor Fault – EPS	Electrical fault in the AC motor driving fuel-system actuation; causes abnormal load/torque behaviour and may disturb downstream fuel delivery.
FM2 – FS Nozzle Switch Open – EPS	Open-circuit or switching fault in the nozzle-control path; disrupts intended nozzle actuation and may alter fuel distribution.
FM3 – FS Valve Switch Open – EPS	Electrical switching fault affecting a fuel-system valve command; causes incorrect valve positioning and associated pressure/flow deviations.
FM4 – Engine Bleed Valve Switch Open – EPS	Electrical/control fault in the engine bleed-valve command path; changes bleed extraction and propagates into ECS pneumatic/thermal behaviour.
FM5 – ECS TCV Switch Open – EPS	Electrical actuation fault in the ECS temperature-control-valve path; causes incorrect pack temperature regulation and strong ECS thermal deviations.
FM6 – AC Lamp Instr. Switch Open – EPS	Instrument-bus switching fault producing local electrical load anomalies and possible disturbance in instrument-supply behaviour.
FM7 – AC Lamp Fluoro Switch Open – EPS	Lighting/load-path switching fault causing local AC bus load redistribution; useful as a predominantly electrical disturbance case.
FM8 – Pump External Leakage – Fuel	External fuel loss from pump or associated interfaces; reduces delivered pressure/flow and may propagate into engine-performance deviations.
FM9 – Pump Internal Leakage – Fuel	Internal recirculation or seal degradation within the pump; lowers useful outlet pressure and delivery efficiency.
FM10 – FOHE Clogging – Fuel	Restriction in the fuel-oil heat exchanger; increases pressure drop and reduces effective fuel throughput.
FM11 – FOHE Leakage – Fuel	Leakage or structural failure in the fuel-oil heat exchanger; causes abnormal pressure loss and degraded fuel/thermal integrity.
FM12 – Fuel Nozzle Clogging – Fuel	Reduced nozzle flow area due to deposits or blockage; distorts fuel metering and can propagate into combustion-related engine signatures.
FM13 – Reduced Pump RPM – Fuel	Under-speed pump operation; lowers delivered fuel pressure and flow, affecting subsystem supply stability.
FM14 – LPT Blade Broken – Engine	Structural damage in the low-pressure turbine; reduces turbine efficiency and alters shaft-work extraction and gas-path behaviour.
FM15 – LPC Fouling – Engine	Surface contamination in the low-pressure compressor; reduces flow capacity/efficiency and shifts pressure–temperature response.
FM16 – HPT Blade Broken – Engine	High-pressure turbine blade damage; alters turbine efficiency, gas temperature distribution, and spool dynamics.

Fault Mode	Compact physical description
FM17 – HPC Contamination – Engine	Degradation in the high-pressure compressor due to deposits/contamination; perturbs discharge conditions and engine thermodynamics.
FM18 – Fan FOD – Engine	Foreign-object damage at the fan; degrades aerodynamic performance and disturbs airflow, pressure, and efficiency.
FM19 – Bleed Valve Angle Fault – Engine	Incorrect bleed-valve position; changes bleed mass flow and directly affects both engine operation and ECS input conditions.
FM20 – CDP Leak – Engine	Leak at or near compressor discharge pressure location; reduces available compressed air and perturbs downstream pneumatic supply.
FM21 – PHX Fouling – ECS	Reduced heat-transfer effectiveness in the primary heat exchanger; leaves downstream air hotter than expected.
FM22 – PHX Cold Mass Flow Blockage – ECS	Restricted ram-air-side cooling through the primary heat exchanger; reduces heat rejection and elevates pack temperatures.
FM23 – SHX Fouling – ECS	Fouling in the secondary heat exchanger; impairs second-stage heat rejection and alters downstream thermal profile.
FM24 – ACM Mechanical Efficiency Loss – ECS	Degradation in air-cycle-machine compressor/turbine work exchange; reduces cooling effectiveness and changes pack pressure/temperature behaviour.
FM25 – RAM Mass Flow Blockage – ECS	Restriction in the ram-air cooling stream; reduces ECS heat-sink capability and increases thermal-chain temperatures.

4.2.2 Data and Methods

This study utilises a high-fidelity Virtual Aircraft Model (VAM), which integrates coupled simulation models of four critical subsystems to capture realistic interactions. The data collected for this work only considers the ground-level runs, and the descriptions of the subsystems are as follows.

The framework starts with pulling master cross-subsystem aircraft dataset stored in `DataB.xlsx`, together with a separate sensor-attribute table, `sensor_attributes.xlsx`, used during the MOSOF optimisation stage. In the public MATLAB workflow, `DataB.xlsx` is the primary input to the NDCI-versus-mRMR benchmarking scripts and to the auxiliary fault-signature/FMECA support analysis, while `sensor_attributes.xlsx` provides non-diagnostic attributes such as cost, reliability/MTBF, and, where available, weight for stakeholder-specific optimisation.

The chapter-level dataset is organised around four aircraft subsystems: Engine**,** Fuel**,** EPS**,** and ECS. In the public analysis scripts, the fault labels are mapped by subsystem as follows: **FM1–FM7** correspond to the electrical subsystem, FM8–FM13 to the fuel subsystem, **FM14–FM20** to the engine subsystem, and FM21–FM25 to the ECS subsystem. The same public workflow uses a severity variable and includes healthy samples together with faulted samples so that subsystem-specific ranking, nested cross-validation, and platform-level optimisation can all be performed from the same master dataset.

The public repository also makes clear that this chapter is based on a high-fidelity virtual aircraft model coupling the engine, fuel, EPS, and ECS, and that the diagnostic benchmarking is performed under a rigorous nested cross-validation protocol before the airline-case MOSOF optimisation is applied. By inspection of the public script definitions, the candidate variables are grouped explicitly by subsystem in code, with a large engine variable set and smaller fuel, electrical, and ECS sets, reflecting the heterogeneous sensing density across the platform. This is important to state in the thesis because the platform-level results are driven by a structured, subsystem-aware dataset rather than by a single undifferentiated feature matrix.

The public codebase contains both a compact support-data convention and the final expanded grouped-variable convention. A legacy support script documents a layout with sensor-reading columns followed by `FaultMode` and `Severity`, whereas the final benchmarking and optimisation scripts define the working variable groups directly in code and map them to the four subsystems used in the chapter. For the purposes of Chapter 4, the thesis should describe the final grouped-variable convention used in the benchmarking and optimisation workflow**,** while noting that all cases retain explicit fault-mode and severity labels.

Engine: The model chosen for the engine digital twin is the Pratt & Whitney JT9D open-source turbofan engine model provided by T-MATS software in MATLAB Simulink [21].

Fuel System (FS): A Simulink-based digital twin of a laboratory fuel rig, capable of simulating pump degradation, blockages, and other flow-related faults.[4]

Electrical Power System (EPS): A simulation model of the EPS, which models generator load dynamics with an Adaptive Neuro-Fuzzy Inference System for its diagnosis (ANFIS) [4]

Environmental Control System (ECS): Proprietary models that simulate bleed air management and thermal performance of heat exchangers [22].

The integrated VAM platform allows for the simulation of cascading faults, providing a rich dataset for training and validating diagnostic algorithms [23]. Raw simulation data is pre-processed to remove constant or nearly constant channels, and the remaining signals are then z-scored. An initial redundancy check removes one sensor from any pair with a Pearson correlation greater than 0.995.

To rigorously compare the performance of NDCI and mRMR, a **nested, aggregated cross-validation** protocol was employed. This method prevents information leakage from the test set into the feature selection process, yielding more reliable and less optimistic performance estimates. An outer loop partitions the data for training and testing. Within each training partition, an inner cross-validation loop is used to determine the optimal number of sensors (k) for a given ranking method. A classifier is then trained on the full training partition using the top k sensors and evaluated on the held-out test set. To account for class imbalance in the fault data, balanced accuracy is used as the primary performance metric.

4.2.3 Classifier Evaluation

To decouple *ranking* effects from *learner* effects and to select a sensible default learner for deployment, an inner cross-validation (CV) search is performed within each outer-fold training split. The classifier catalogue comprises class-weighted Bagged Trees (Bag), SVM with RBF kernel (svmRBF), subspace k-nearest neighbours (subKNN), and Kernel Naïve Bayes (nbKernel); inverse-frequency misclassification costs address class imbalance for all learners. For a given ranking method and subsystem, features are added stepwise in the learned rank order, inner-CV balanced accuracy is recorded across steps, and a one-standard-error rule selects the number of k features. The learner is chosen by the inner-CV peak under the same rule; all steps use training partitions only to prevent information leakage into the outer test. The best performers in each subsystem for detection and isolation tasks are shown in Table 4-2.

Table 4-2 Best-performed classifier comparison per subsystem

Subsystem	Task	Best classifier
Engine	Detection	Bag
Engine	Isolation	Bag
Fuel	Detection	Bag
Fuel	Isolation	nbKernel
EPS (Elec)	Detection	Bag
EPS (Elec)	Isolation	svmRBF
ECS	Detection	Bag
ECS	Isolation	Bag

Class-weighted Bagged Trees are best suited for detection tasks across all subsystems. For isolation, Bag remains the preferred learner on Engine and ECS, **SVM-RBF** is preferred on EPS, and **nbKernel** is preferred on Fuel. These choices are consistent with subsystem-specific decision-boundary geometry and the diversity of available sensors; they are used as the default learners when reporting deployed confusion matrices and quantitative analysis.

4.3 Results

This section reports a two-step path from a platform-level recommendation to a subsystem run for comparison:

Platform symptom vectors.
How each fault perturbs sensors across the entire aircraft is visualised by plotting normalised deviations from the healthy baseline. These platform-level plots show where a fault's influence is concentrated and how it propagates into

other subsystems (e.g., pneumatic → thermal → electrical). They justify the subsequent design choices by revealing genuine cross-subsystem couplings, while also showing that a single platform-wide ranking is not robust for this dataset.

Subsystem-level, quantitatively comparable rankings and classifiers. Because a global ranking was unstable (dominated by a few high-variance channels and heterogeneous scaling), feature rankings and classifiers are evaluated per subsystem. This isolation method prevents cross-subsystem leakage during training and yields comparable, leakage-free performance estimates for the Engine, Fuel, ECS, and EPS subsystems. The resulting candidate sensors and their measured diagnostic value are then handed to MOSOF (Section 4), which composes platform suites that mix sensors across subsystems under cost/reliability constraints.

In short, the figures at the start of this section present platform-level symptom vectors for interpretation. Section 4 returns to the platform to optimise and select a cross-subsystem suite.

4.3.1 Cross-Subsystem Synergies and Feature Ranking

The platform-level “symptom vectors” that visualise how each fault mode perturbs sensors across subsystems are shown before presenting the quantitative, subsystem-level rankings. For a given fault and severity, each sensor reading is expressed as a normalised deviation from the healthy baseline; plotting these deviations as a vector reveals where the effect concentrates and how far it propagates into other subsystems.

Directly optimising a single platform-wide ranking with either NDCI or mRMR proved unstable and physically not sensible for this dataset (dominated by a few high-variance channels and heterogeneous scaling across subsystems). Therefore, the platform symptom vectors (PSV) are used for interpretation only and perform the quantitative comparisons per subsystem. The following figures establish the cross-subsystem context that motivates that design choice.

a-)

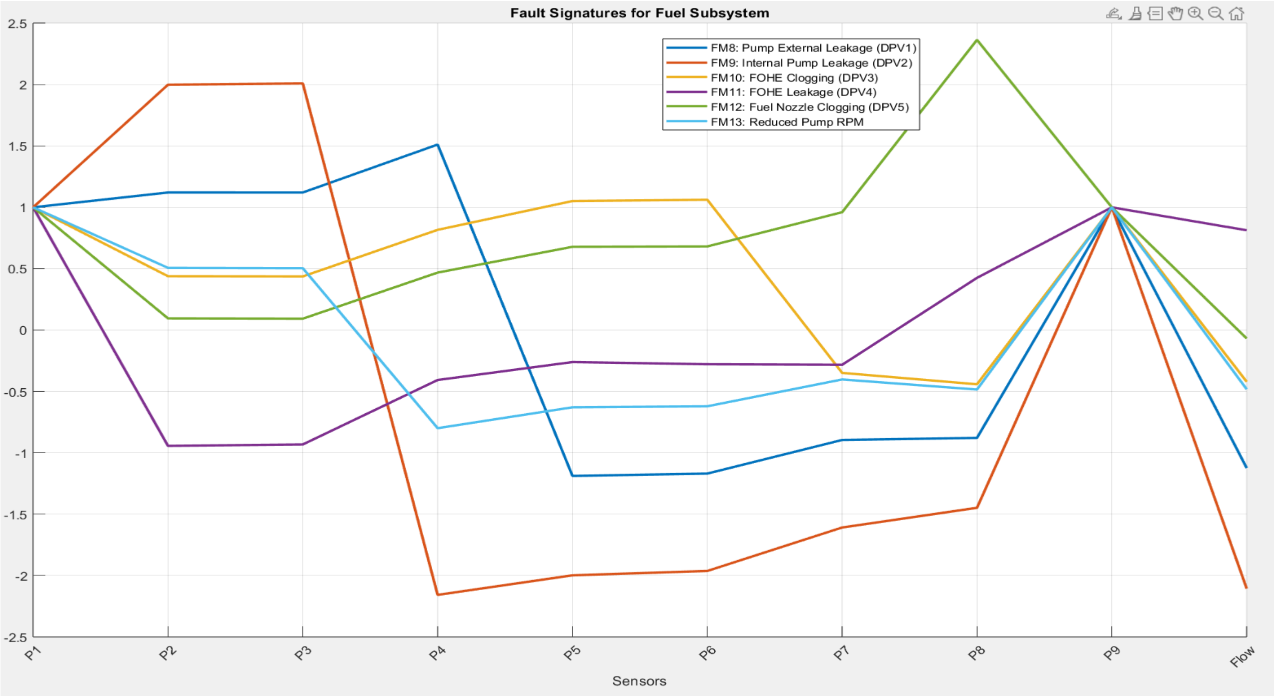

b-)

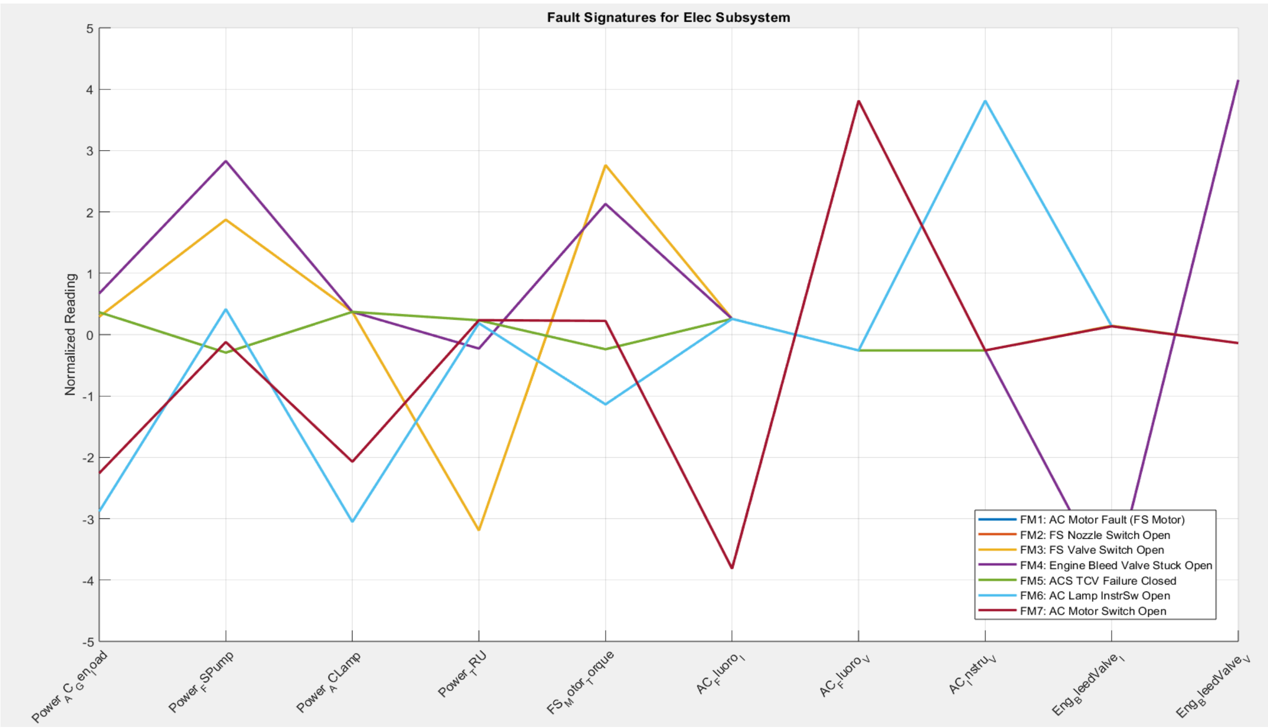

Figure 4-2 a) Fuel fault signatures and b) EPS fault signatures— example of system symptom vectors

Figure 4-2a, the six Fuel faults (FM8–FM13) produce distinct signatures over the pump pressures (P1–P7) and the Flow channel. Pump-leakage and reduced-RPM faults show broad, moderate deviations across pressures, whereas FOHE-related faults concentrate around the manifold taps.

In Figure 4-2b, EPS faults (FM1–FM7) generate characteristic responses on power-related channels (e.g., generator load, lamp/instrument buses). Notably, faults that involve bleed air actuation (e.g., valve/motor issues) imprint on Electrical measurements via load changes, illustrating the coupling between electrical actuation and the pneumatic path.

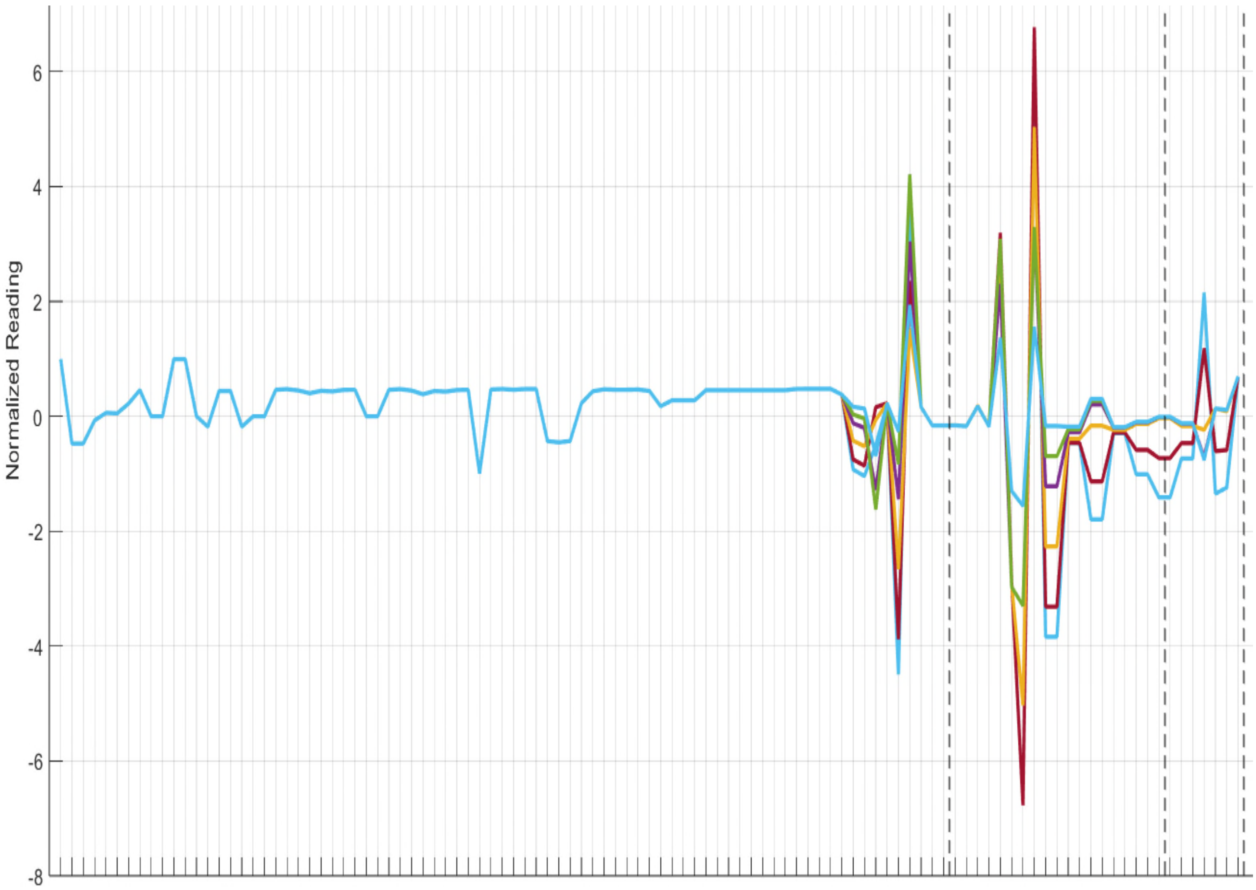

Figure 4-3 FM 5 severity sweep (0.1 incremental) across PSV

FM5, represented in Figure 4-3, is an EPS fault named “ECS Temperature Control Valve Switch Open”. Curves show the normalised deviation across all sensors as the FM5 severity increases; vertical dashed lines mark subsystem boundaries. Excursions of the FM5 EPS fault grow non-linear in the ECS chain,

and other electrical load channels exhibit similar responses. These patterns indicate how the severity of one fault mode affects the sensor readings of different subsystems. The informative content is highly localised to the Fuel subsystem, with limited propagation to other subsystems. The concentration of severity-dependent changes around these regions is the principal cross-subsystem synergy leveraged by MOSOF later.
Where the “macro effects” appear and how they shape the solution are derived from the PSV of each fault mode. The platform-level symptom-vector figures establish the macro pathways highlighted in the introduction, by originating each PSV of each FM:

Fuel → Engine. Fuel system faults modulate Flow and manifold pressures, and through combustion, alter engine temperature/enthalpy signatures. This is visible in the Fuel and Engine symptom vector plots and reappears in the final suite via Flow (Fuel) and engine thermodynamic channels, with high NDCI scores.

Engine bleed-air → ECS thermal chain. Electrical FM5 (ACS TCV open) drives strong excursions along the ECS heat exchanger temperature chain and the engine bleed/mass flow variables; the cross-subsystem effect distribution shows the largest bars on ECS thermal sensors, with a concurrent spike in bleed air mass flow. NDCI prioritises these same regions and persists in multi-objective selection.

Pneumatic/thermal → electrical load. Actuation and valve states couple back into electrical channels through load changes, explaining why some EPS power measurements enter the Pareto-efficient suites despite modest stand-alone separation.

These platform-level visuals for each FM’s effects on the sensor readings of other systems and their evaluation are based on the system interactions. The Platform-wide rankings did not yield sensible results for this dataset; however, the cross-subsystem effects are real and can be best exploited by building subsystem-level rankings and then allowing MOSOF to combine sensors across subsystems. The evaluation of the stepwise accuracy for platform-level results is shown in Figure 4-4

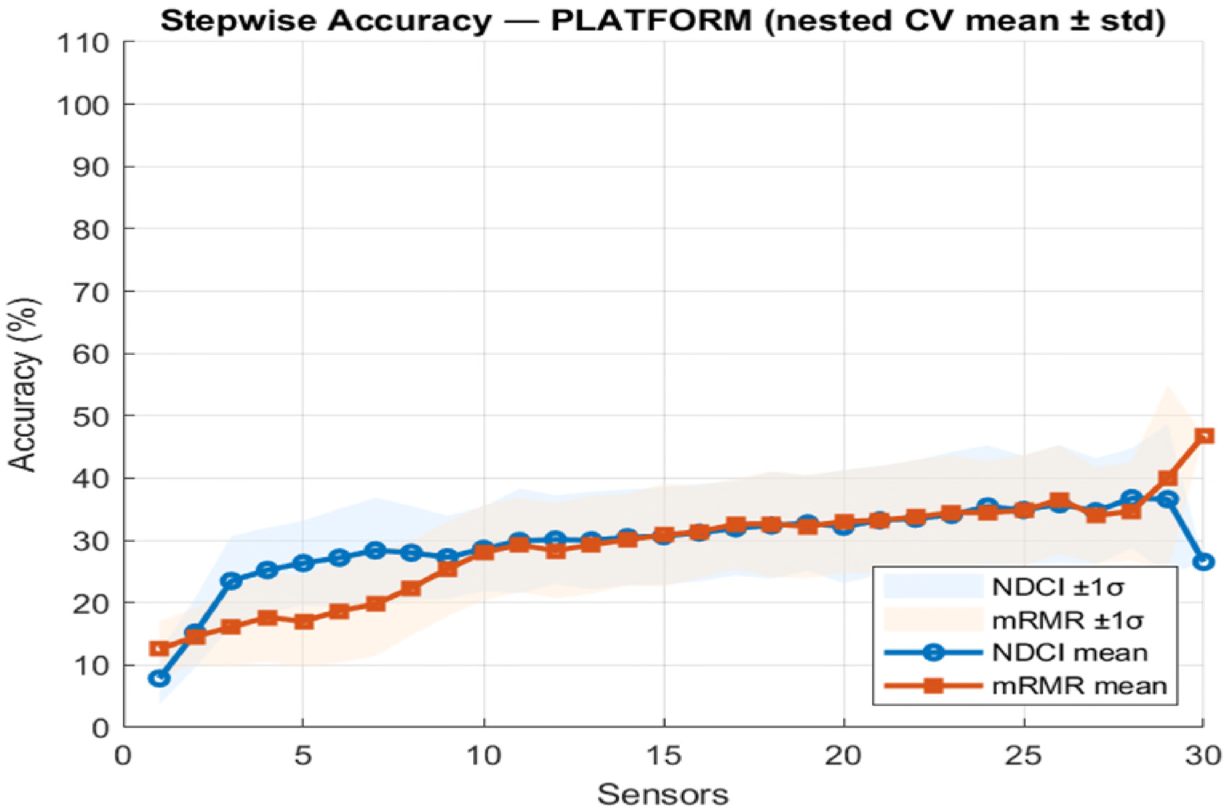

Figure 4-4 Stepwise accuracy for the platform-level evaluation

To test whether a single platform-wide ranking could be deployed directly, a global experiment pooled all candidate sensors across subsystems. It evaluated NDCI and mRMR by adding features stepwise based on the technique's ranking. The platform-level stepwise accuracy in Figure 4-4 shows a low early rise (≈ 20–25% by 3–5 sensors) followed by a broad plateau around 30–35% with wide fold-to-fold variability; additional sensors beyond ≈ 10 sensors yield only marginal accuracy gains. This behaviour indicates that a single global ranking is not yielding reasonable results for this dataset; the heterogeneous scales and class distributions across subsystems diminish discriminability when features are mixed globally, yet the cross-subsystem effects are real, as established by the symptom-vector visuals above. Consequently, all quantitative ranking comparisons are performed per subsystem (Engine, Fuel, EPS, ECS) to obtain stable, leakage-free estimates.

4.3.2 Baseline vs. Nested Cross-Validation Performance

Baseline experiments were conducted using a simple five-fold cross-validation protocol on a dataset with limited samples, which included three degraded, faulty readings with varying severity levels, as well as a healthy reading for each fault mode. The result of the baseline runs yielded an optimistic outcome; therefore, a more comprehensive technique, nested cross-validation, is applied to enhance the correctness of the results. The nested cross-validation protocol:

Goal: Obtain unbiased performance estimates while comparing ranking methods fairly.

Outer loop (testing): Data are partitioned into outer training and test folds (stratified by fault mode). The test fold remains untouched until the final evaluation in that outer iteration.

Train-only ranking: Within each outer training set, NDCI and the baselines (e.g., mRMR) are computed using training data only. Redundant sensors are pruned on training faults.

Inner loop (model selection): On the training faults only, an inner K-fold loop evaluates a single, fixed classifier while adding sensors stepwise in the given rank order (top-1, top-2, …). At each step, the inner-CV mean ± std accuracy is recorded. The number of sensors, k, is chosen by a one-standard-error (SE) rule from this inner loop.

Fit and test: Using the top-k sensors for that ranking, the whole outer training set is retrained on and predicts the held-out outer test faults.

Repeats and aggregation: The entire process is repeated with different partitions; outer-test predictions are pooled to form the confusion matrices and summary metrics (balanced accuracy for detection/isolation).

What each figure means.
• Stepwise curves plot inner-CV mean ± std accuracy vs. the number of sensors added—this reflects the *sample efficiency* of each ranking.
• Tables/summary bars report outer-test performance aggregated over

repeats/folds—this reflects *generalisation*.
• Confusion matrices aggregate outer-test predictions; this reveals which faults are confused in practice.

This protocol ensures that feature ranking, selection of k, and final testing are cleanly separated. All ranking methods are evaluated using the same classifier (class-weighted bagged decision tree ensemble), metrics, and splits, so differences arise from the ranking itself rather than from modelling choices.

As shown in Table 4-3, under these conditions, NDCI demonstrated a significant advantage over mRMR. For the Engine subsystem, NDCI achieved 92.6% balanced accuracy with only 9 sensors, compared to 77.8% for mRMR with 12 sensors. Similar gains were observed for the EPS (73.1% vs. 34.6%) and Fuel (66.7% vs. 58.3%) subsystems.

Table 4-3 Baseline cross-validation results for each subsystem

Subsystem	Method	Balanced Accuracy	Sensors Used
Engine	NDCI	0.926	9
Engine	mRMR	0.778	12
Fuel	NDCI	0.667	5
Fuel	mRMR	0.583	7
EPS	NDCI	0.731	8
EPS	mRMR	0.346	12
ECS	NDCI	0.812	6
ECS	mRMR	0.812	8

To obtain more reliable performance estimates, switched to a nested, aggregated cross-validation protocol with an increased sample size of around 40 readings per fault class. The results, summarised in Table 4-4, show an expected decrease in overall accuracy scores due to the more challenging and realistic validation scheme. However, the performance advantage of NDCI was largely preserved. For the Engine, NDCI's accuracy remained high at 88.6%, while mRMR's performance dropped sharply to 69.0%. For the ECS, NDCI achieved 67.7% accuracy, compared to 52.0% for mRMR. The Fuel subsystem remained a challenge for both methods, with mRMR showing a slight advantage (53.0% vs. 48.7%), an outcome attributed to the limited diversity of available sensor types for that subsystem.

Table 4-4 Nested, aggregated cross-validation results

Subsystem	Method	Balanced Accuracy	Sensors Used
Engine	NDCI	0.886	10
Engine	mRMR	0.690	12
Fuel	NDCI	0.487	5
Fuel	mRMR	0.530	5
EPS	NDCI	0.518	7
EPS	mRMR	0.510	8
ECS	NDCI	0.677	6
ECS	mRMR	0.520	7

The large gap between the simple baseline and the nested protocol is methodological. The baseline used a single pass cross-validation in which feature ranking and the choice of the number of sensors were not separated from evaluation; this allows information from the test fold to influence selection (optimistic bias under small samples). The nested protocol enforces separation and adds repeatability:

Repeats/folds: 10 repeats; 5 outer folds for testing; 3 inner folds for model selection.

Train/test split: within each outer fold, feature ranking uses training partitions only; held-out tests are never seen until the final evaluation.

What is ranked: training-fault rows with a healthy baseline; redundant sensors pruned at $|\rho| > 0.995$ on training data.

Learner: a class-weighted bagged decision-tree ensemble (identical across methods) with inverse-frequency class costs.

Model selection: inner-CV stepwise curves determine k using a one-SE rule; only the top-k sensors are carried to the outer test.

Metric alignment: stepwise panels show inner-CV accuracy (mean ± sd); outer-test summaries report balanced accuracy aggregated across repeats.

Under this protocol, the performance drops from the inflated baseline to values that agree with physical intuition and platform-level symptom vectors. The lesson is straightforward: without nesting, the evaluation overstates what a deployable system would achieve; with nesting, estimation variance tightens, and conclusions become defensible.

The aggregated confusion matrices from the nested protocol provide a detailed view of diagnostic performance. Figure 4-5 shows the confusion matrix for the EPS subsystem, and Figure 4-6 shows the matrix for the Engine subsystem, both using the NDCI-selected suites.

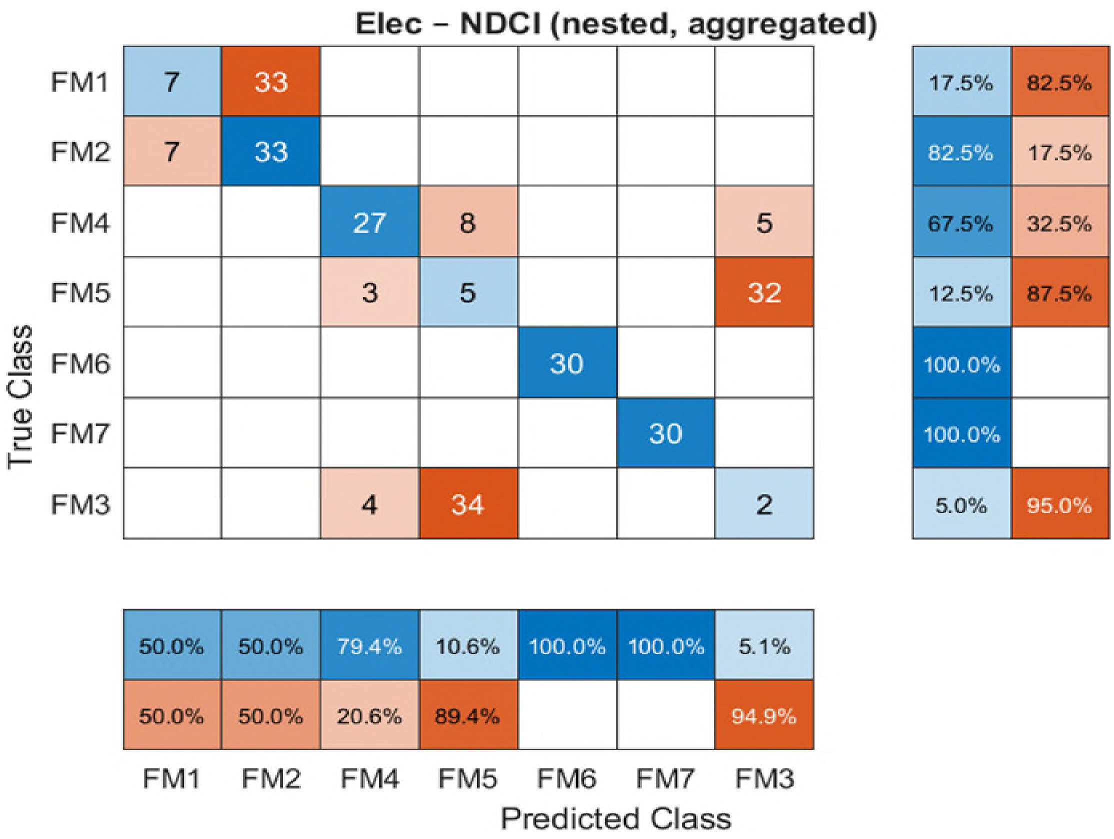

Figure 4-5 Confusion matrix for the EPS subsystem using the NDCI-selected suite (nested, aggregated, classifier bagged decision tree)

For the EPS subsystem, unlike the Engine, there are several areas of significant confusion. For example, both FM1 and FM2 are frequently classified as FM2 (33 instances each). Similarly, FM3 is often misclassified as FM5 (34 instances). This visualisation is crucial as it reveals not just the overall accuracy but also the specific failure modes that the current sensor suite struggles to disambiguate.

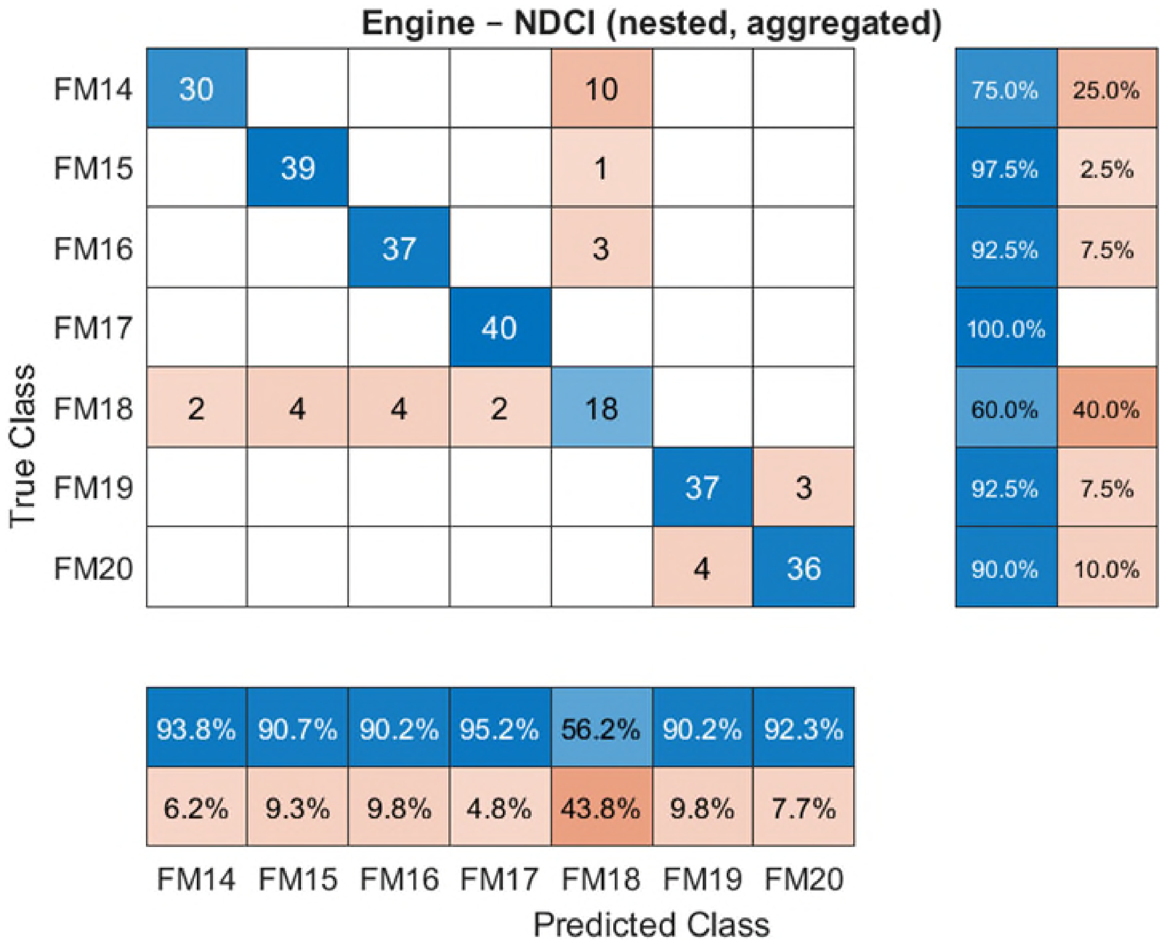

Figure 4-6 Confusion matrix for the Engine subsystem using the NDCI-selected suite (nested, aggregated)

The Engine confusion matrix visualises the classification performance using the NDCI-selected sensors. Each row represents the true fault class, and each column represents the predicted class. The strong diagonal, indicated by the dark blue cells with high counts (e.g., 39 for FM15, 40 for FM17), demonstrates a high true positive rate. Off-diagonal cells in orange indicate misclassifications, which are relatively infrequent, confirming the high overall accuracy reported in Table 4-4.

A repeated nested cross-validation pipeline is used for each subsystem independently (Engine, Fuel, ECS, EPS). In every outer fold, feature rankings are learned on the training partition only, and an inner K-fold loop evaluates a single, fixed classifier while adding sensors stepwise in rank order (top-1, top-2, …). The classifier is a class-weighted bagged decision-tree ensemble (weights inversely proportional to class frequency) and is kept identical for all ranking methods to ensure a fair comparison. At each step, the inner CV mean ± standard deviation of accuracy forms the stepwise curves; the selected number of sensors, k, for each ranking is chosen by a one-standard-error (SE) rule. The model is then retrained on the outer-fold training data with the top k sensors and evaluated on the held-out test faults.

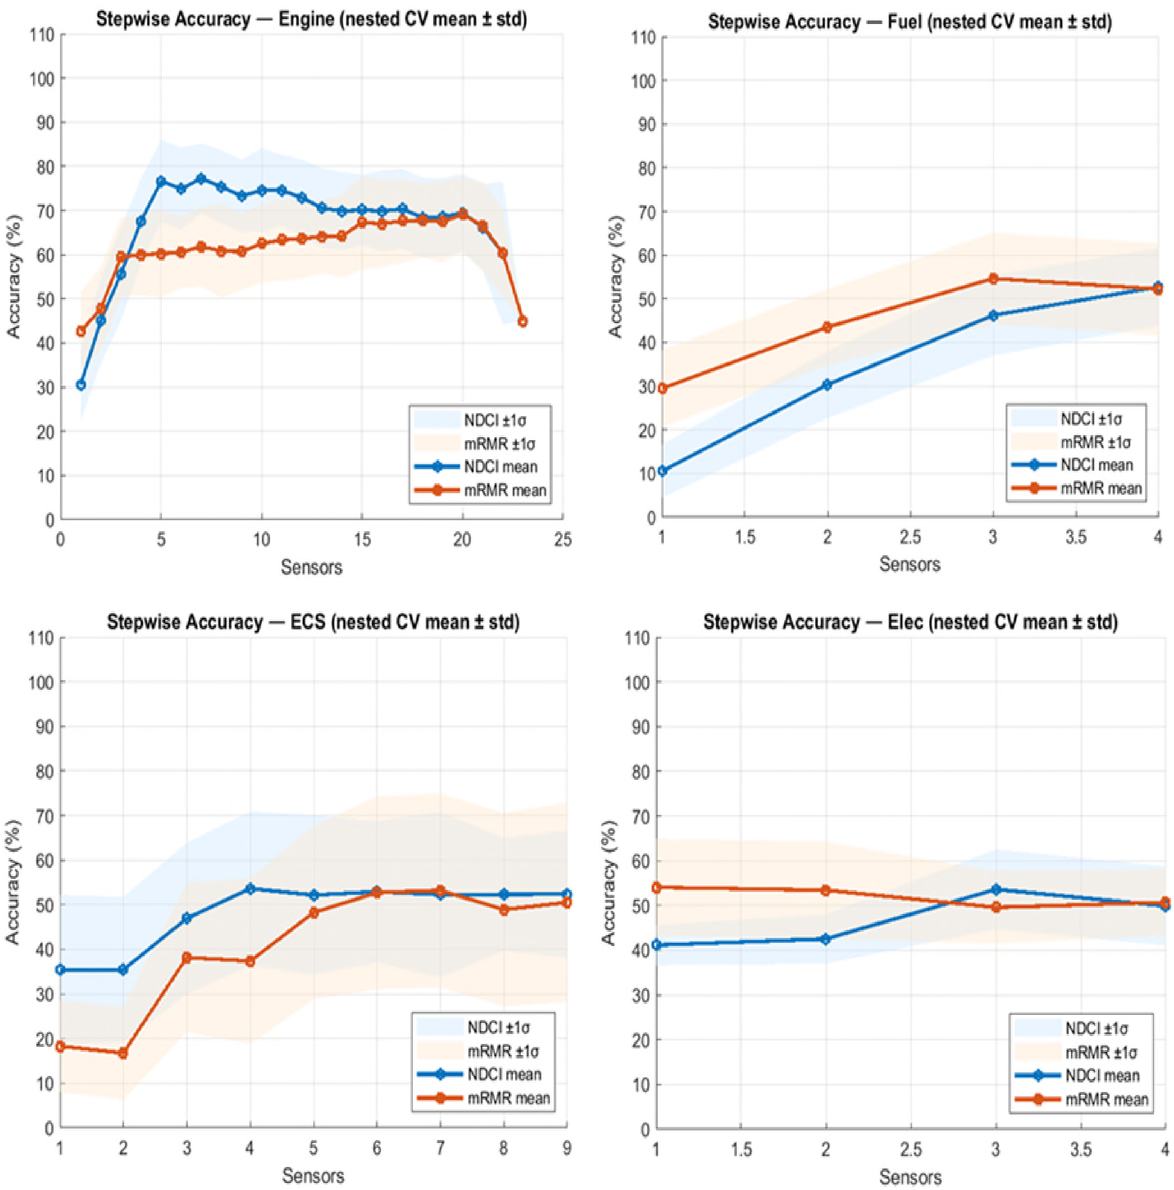

Figure 4-7 Stepwise accuracy for the Engine subsystem (nested CV mean ± std): NDCI vs. mRMR

Figure 4-7 consolidates the stepwise mean ± std curves for all four subsystems, comparing NDCI to mRMR under the same nested protocol. The Engine and ECS panels indicate that NDCI achieves higher accuracy with fewer sensors, thereby improving sample efficiency. On the Engine, accuracy rises rapidly and approaches ~75% by about five sensors with NDCI, whereas mRMR requires more sensors to achieve a similar level. For the Fuel subsystem, mRMR holds a slight advantage; consistent with the smaller and less diverse Fuel sensor set, where NDCI's uniqueness term becomes more restrictive. The EPS panel is more competitive, with NDCI peaking slightly higher around three sensors.

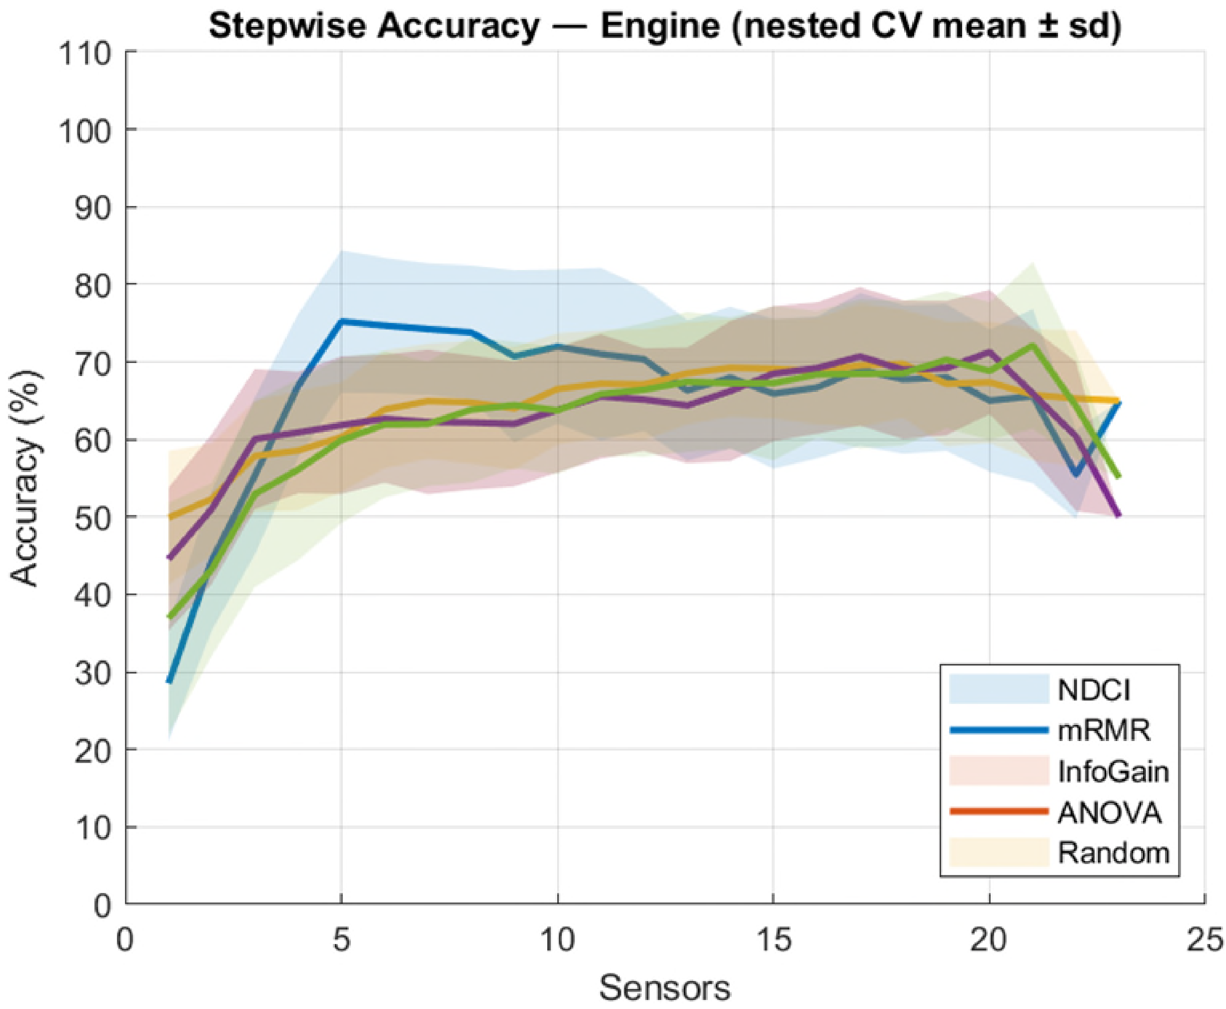

Figure 4-8 Stepwise accuracy for the Engine subsystem (nested CV mean ± std): NDCI, mRMR, Info-Gain, Anova, and Random

To probe method sensitivity on a representative case, Figure 4-8 adds two additional ranking baselines on the Engine subsystem. Info-Gain (mutual information) and ANOVA-F, together with a Random ranking obtained by randomly permuting the candidate list. Info-Gain and ANOVA closely track the mRMR family, while Random grows slowly and plateaus early, confirming that the gains stem from the ranking rather than the stepwise evaluation procedure. Across the range, NDCI remains the most sample-efficient, achieving high accuracy with a compact subset by jointly rewarding separation power, severity sensitivity, and uniqueness in ranking.

All stepwise plots in this subsection are per subsystem (the model considers Healthy plus that subsystem's fault modes only) to isolate the effect of the ranking method and avoid cross-subsystem leakage. The cross-subsystem value is then exploited in the platform-level MOSOF study, where the final suite combines sensors from multiple subsystems.

To understand the underlying reasons for the different sensor rankings, figures 4-9 to 4-12 compare the NDCI component breakdown with the mRMR rank order for each subsystem. Each figure has two panels for the same subsystem:

Left panel — NDCI components (stacked bars).
Each horizontal bar = one sensor. The bar is divided into three normalised NDCI components: SP (Separation Power), S (Severity Sensitivity), and U (Uniqueness). The total bar length is proportional to the composite NDCI (SP+S+U, i.e., higher is better).
Ordering: sensors are sorted by composite NDCI, so "better" sensors appear lower in the plot. In other words, read the left panel from bottom to top: the bottom bars represent the highest NDCI sensors.

Right panel — mRMR rank values (shorter is better).
This panel shows only the rank order produced by the mRMR baseline. The bar length equals the numerical rank (1 = best, 2 = second-best, …).
Important: the y-axis lists sensors alphabetically, not by rank. Consequently, rank 2 is not necessarily underneath rank 1. To interpret the panel, look at the number along the x-axis (bar length): shorter bars = better ranks.

This side-by-side view compares the two methods for ranked sensors: NDCI prioritises sensors with strong and balanced SP/S/U contributions. At the same time, mRMR focuses on relevance with redundancy control, which can favour different channels.

For example, in the Engine (Fig. 9), NDCI elevates sensors such as W_AfBleed and h_S4 because they exhibit clear fault separation, monotonic severity response, and low redundancy. In contrast, mRMR's top-ranked items (short bars at the right) may cluster around a smaller set of correlated variables that are highly relevant but less unique.

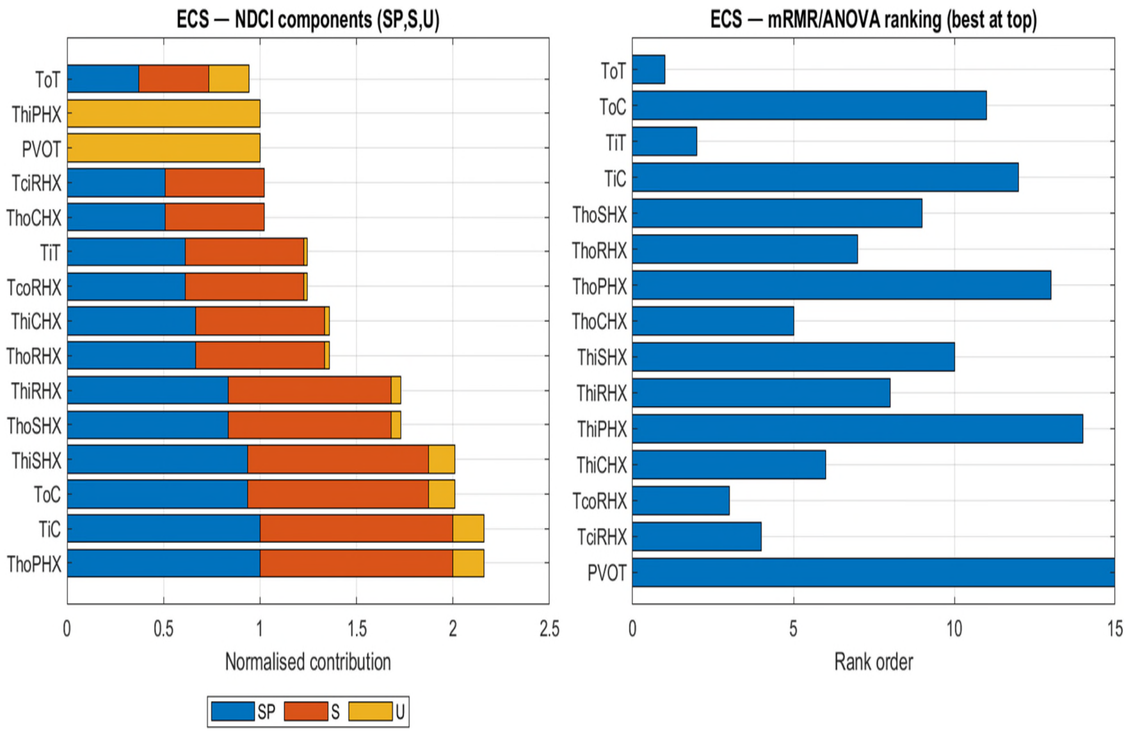

Figure 4-9 Ranking comparison for the ECS subsystem: NDCI components (SP, S, U) vs. mRMR ranking

Figure 4-9 provides a side-by-side comparison of sensor rankings for the ECS. The left panel shows the NDCI score for each sensor, broken down into its three components (SP, S, U). The right panel shows the rank order produced by mRMR. This visual comparison helps to explain why the methods select different sensors, highlighting NDCI's emphasis on a balanced diagnostic contribution.

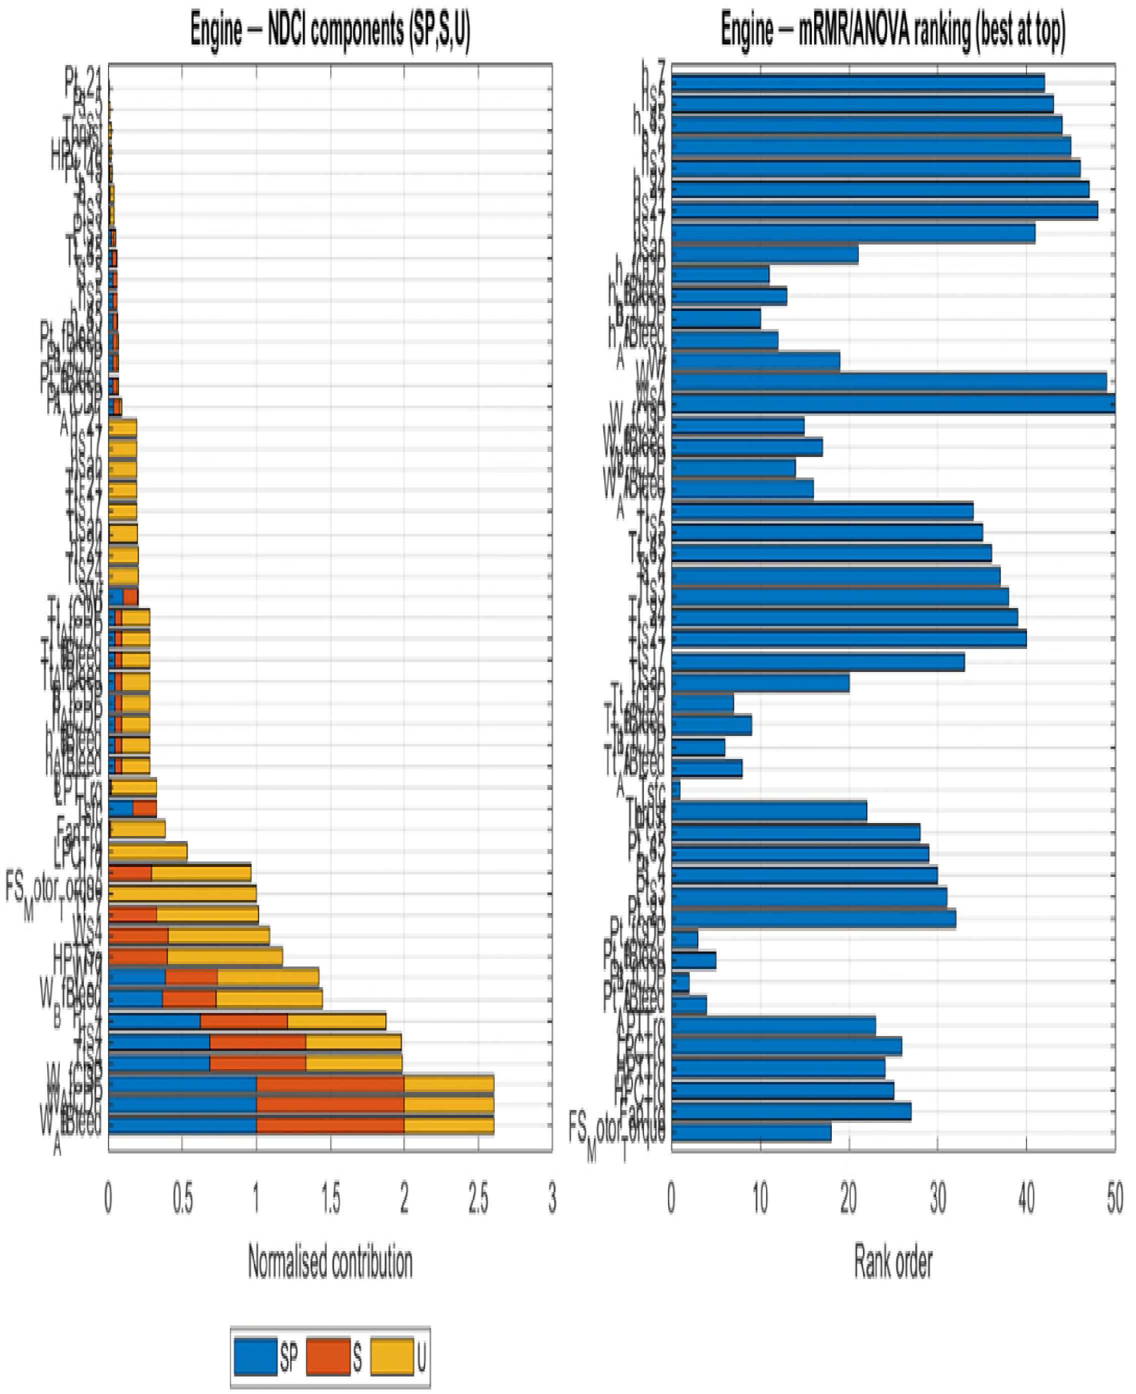

Figure 4-10 Ranking comparison for the Engine subsystem: NDCI components (SP, S, U) vs. mRMR ranking

Similar to the previous figure, this plot (Figure 4-10) for the Engine subsystem reveals the differing priorities of the two ranking methods. The NDCI ranks the sensor readings based on its defined metric, whereas the mRMR ranking on the right is based on label relevance (mutual information) with redundancy control.

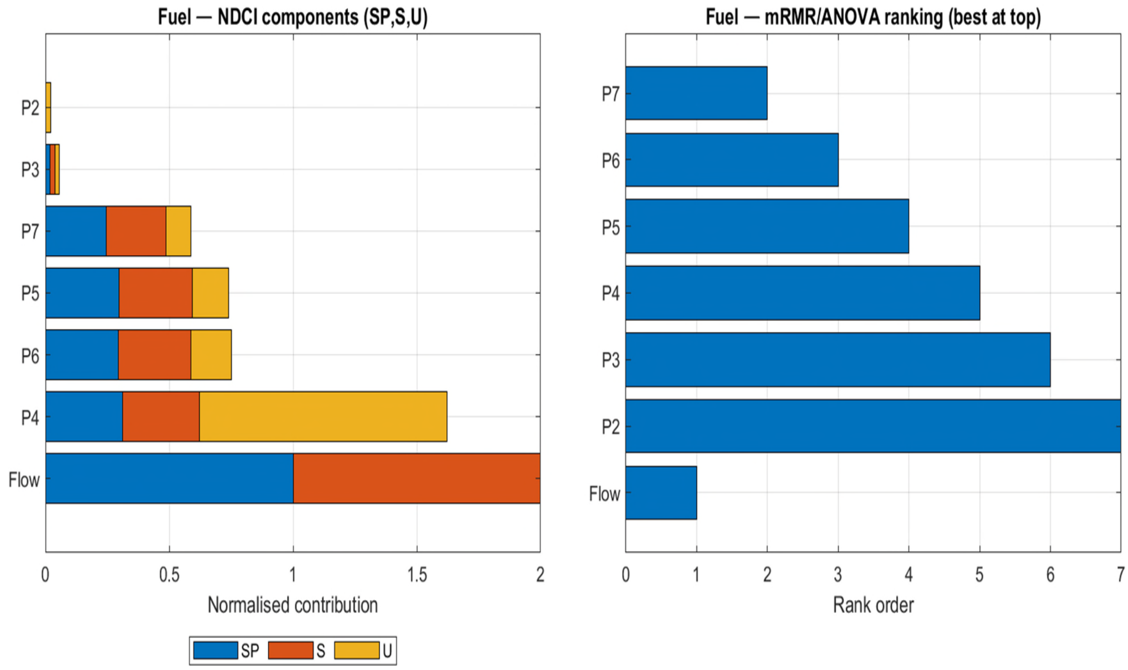

Figure 4-11 Ranking comparison for the Fuel subsystem

For the Fuel subsystem, Figure 4-11 illustrates why NDCI struggled relative to mRMR. The NDCI component plot on the left shows that while Flow has high SP and S, the other sensors (P2-P7) have relatively low scores, especially when uniqueness is considered. This visualisation supports the conclusion that a lack of sensor diversity limited NDCI's effectiveness in this specific case.

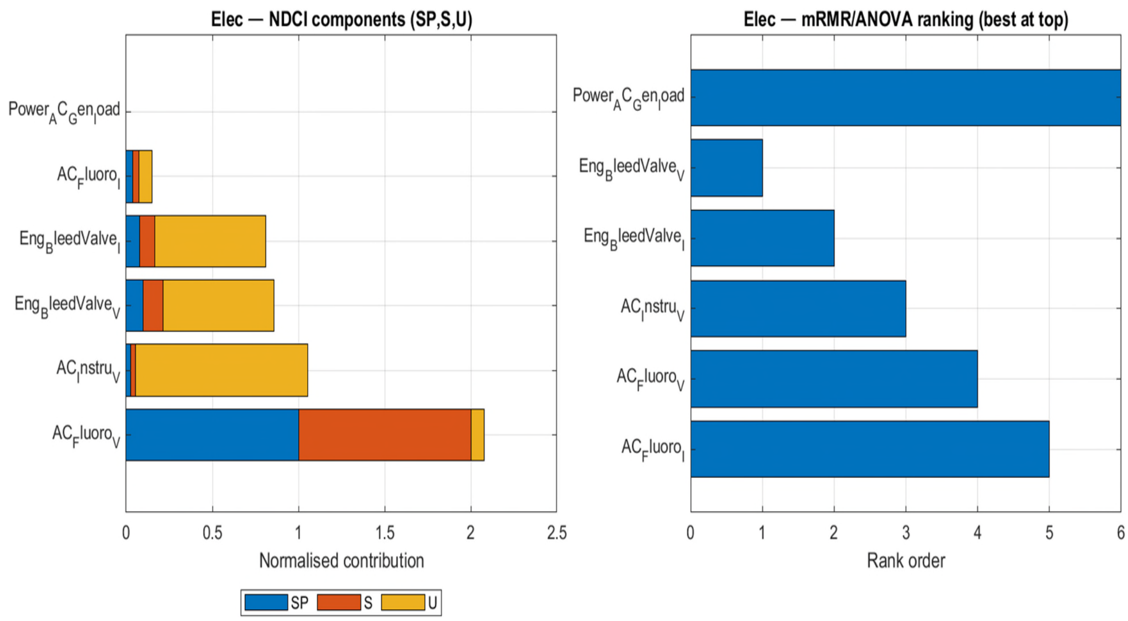

Figure 4-12 Ranking comparison for the EPS subsystem

Figure 4-12 compares the sensor rankings for the EPS. It shows how NDCI prioritises sensors like AC_Fluoro_V and AC_Instru_V, which have high Uniqueness (yellow bar), meaning they provide information that is distinct from other available signals, a key factor for effective fault isolation in the EPS.

The final accuracies for the NDCI-selected suites across all subsystems are visualised in Figure 4-13.

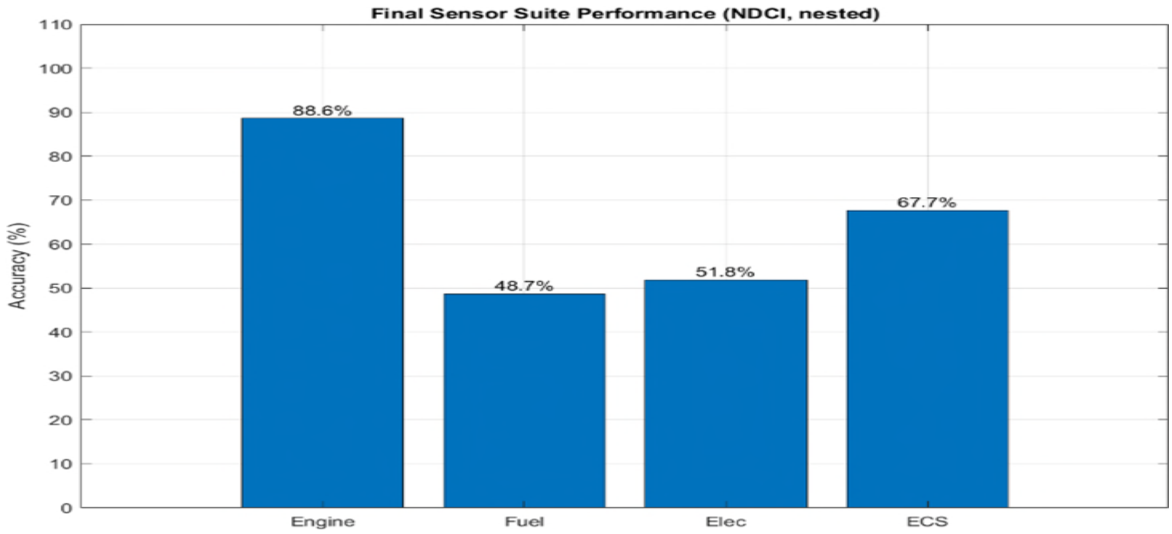

Figure 4-13 Final sensor-suite accuracy comparison (best-k, nested)

The bar chart in Figure 4-13 summarises the peak diagnostic performance (balanced accuracy) achieved by the NDCI-selected sensor suites for each subsystem under the robust nested validation protocol. The Engine subsystem clearly shows the highest diagnostic accuracy at 88.6%, followed by the ECS at 67.7%. The lower performance for the EPS (51.8%) and Fuel (48.7%) subsystems indicates a greater intrinsic difficulty in diagnosing their respective fault modes with the available instrumentation.

For completeness, the confusion matrices for the cases where mRMR competes with (Figures 4-14) and marginally outperforms NDCI (Figure 4-15) are presented.

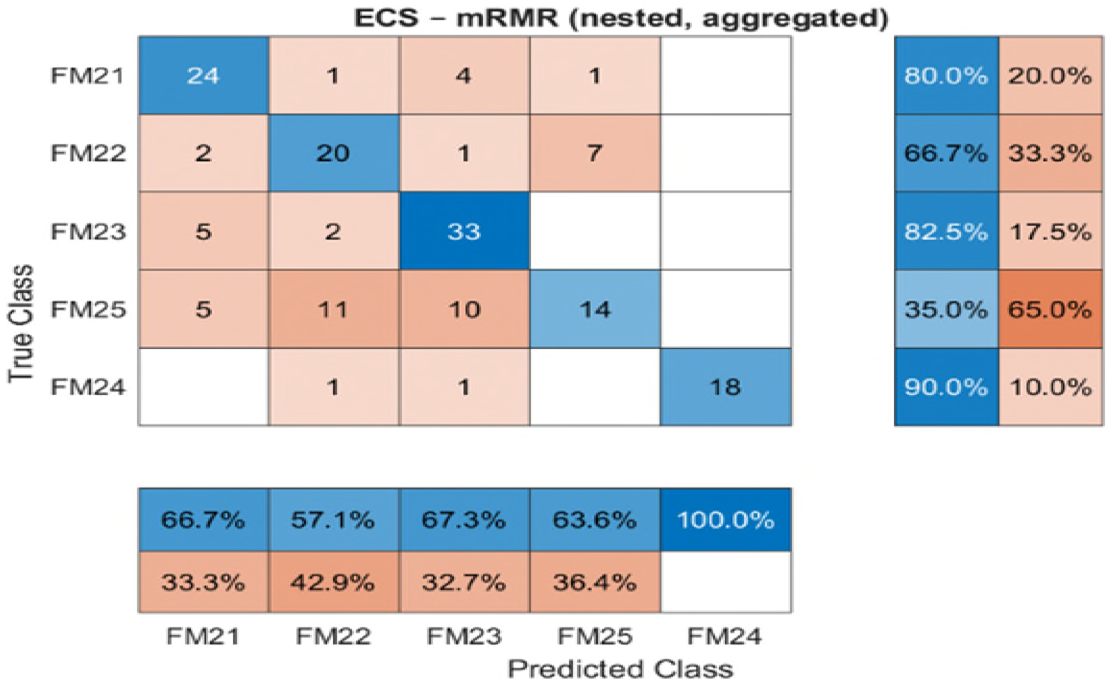

Figure 4-14 ECS confusion matrix for the mRMR suite (nested, aggregated)

Figure 4-14 shows the performance of the mRMR-selected sensor suite for the ECS. While overall accuracy is lower than NDCI's, this view allows for a direct comparison of which specific fault modes mRMR handles effectively.

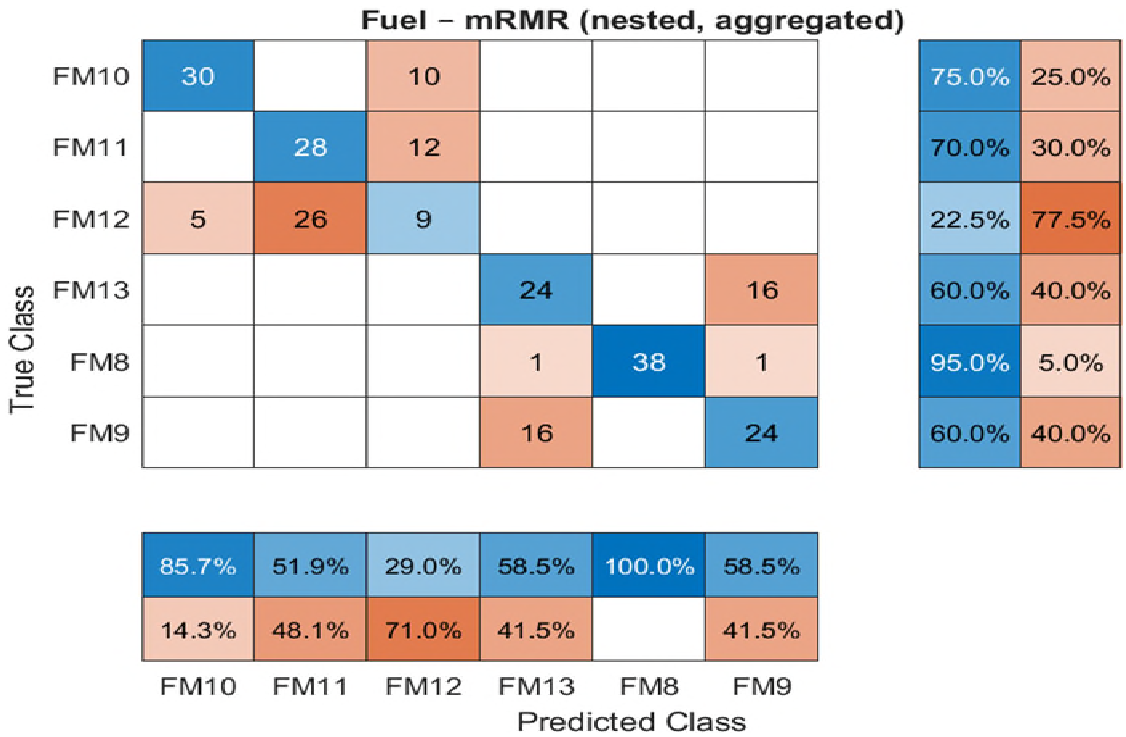

Figure 4-15 Fuel confusion matrix for the mRMR suite (nested, aggregated)

This confusion matrix in Figure 4-15 details the performance of the mRMR-selected suite for the Fuel subsystem, where it achieved a slightly higher balanced accuracy than NDCI. It highlights the specific classification strengths and weaknesses of the mRMR approach in a sensor-constrained environment.

4.4 Airline-Centric MOSOF Trade-off Study

The MOSOF was executed to find a globally optimised sensor suite for an airline stakeholder whose objectives are discussed and evaluated in the previous paper on the ECS case. The resulting Pareto front, containing all non-dominated solutions, is visualised in Figures 4-15 and 4-16. These plots illustrate the trade-off surface between performance, cost, and reliability.

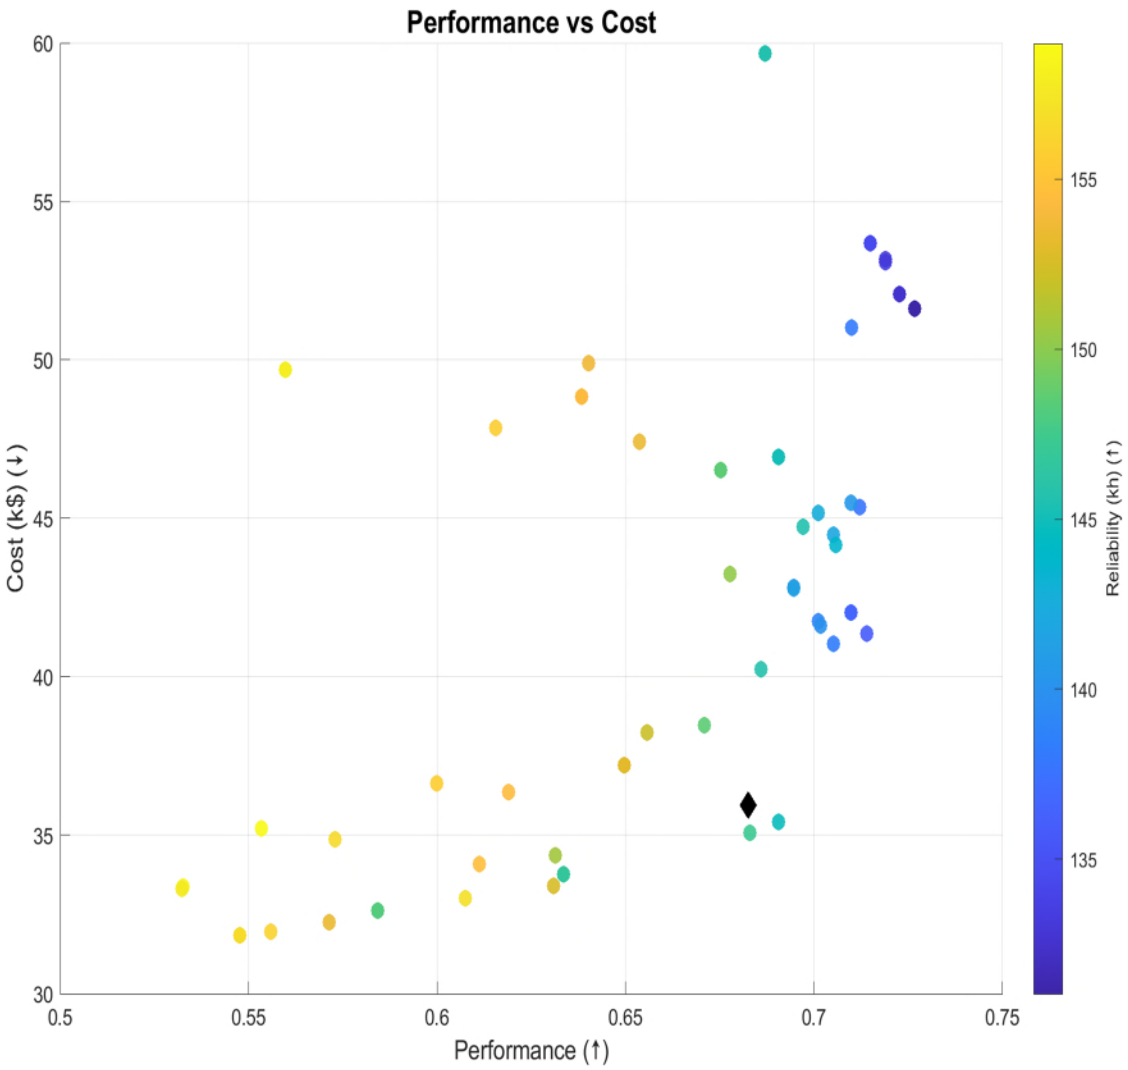

Figure 4-16 Feasible Pareto set (performance vs. cost; colour = reliability)

Figure 4-16 shows the projection of the tri-objective Pareto set onto Performance vs. Cost, with Reliability encoded by colour. Because this is a projection, it does not imply a monotonic relation between cost and performance. In fact, the cloud exhibits three clear regimes:

- a flat, budget floor around $32–36k, where additional spend yields little performance change.
- a moderate-slope region up to ≈0.68 performance; and
- a steep, diminishing-returns region beyond ≈0.69 where sizeable cost increases buy only small gains.

At ≈approximately 0.70 performance, several designs span $40–$ 55,000 with similar performance. This is expected: these points differ along the reliability dimension (colour) and in their sensor mix. Some higher-cost suites utilise pricier channels or incorporate redundancy that enhances or maintains reliability, while others compromise reliability for marginal performance improvements. The black diamond marks the knee—the solution on the front with the most significant normalised gain in performance (and reliability) per unit cost relative to its immediate neighbours (the "elbow" of the curve in this 2D view).

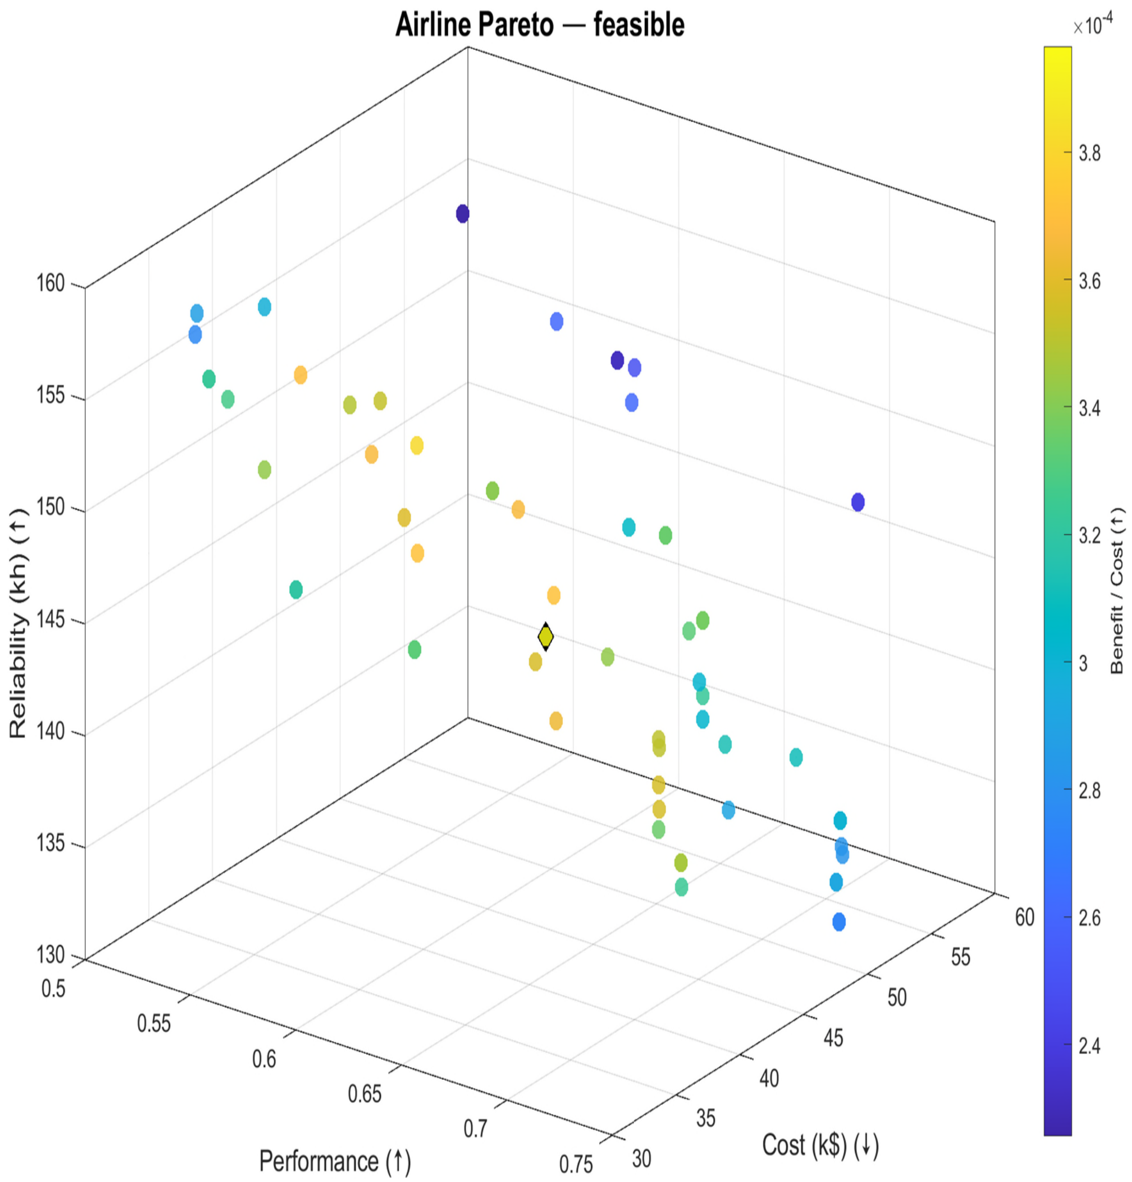

Figure 4-17 3D Pareto front with benefit-to-cost colouring; the knee solution is marked

Figure 4-17 plots the same non-dominated set in three dimensions: Performance (x ↑), Cost (y ↓), and Reliability (z ↑). This view adds two things that the 2D projection cannot show:

Why do points with similar Performance–Cost differ? The third axis reveals their reliability separation. For example, the designs clustered near 0.70 performance in Fig. 16 occupy different heights in Fig. 17 (corresponding to different reliability levels), which explains the wide cost band seen in the projection.

Where the efficient “ridge” lies in 3D, the front forms a curved surface; moving toward higher performance can either lift reliability (good) or flatten/drop it (undesirable). Seeing this surface helps identify regions where small cost increases improve both performance and reliability, versus areas where one improves at the expense of the other.

The colour in Fig. 17 shows a benefit-to-cost ratio (a normalised combination of performance and reliability divided by cost), purely for visual emphasis. Brighter colours flag suites that deliver more capability per dollar; they are not used to define Pareto optimality. The yellow diamond represents the knee solution (same design as the black diamond in Fig. 16), chosen at the highest-curvature region of the front after objective normalisation; it is the point beyond which the incremental cost yields progressively smaller joint gains.

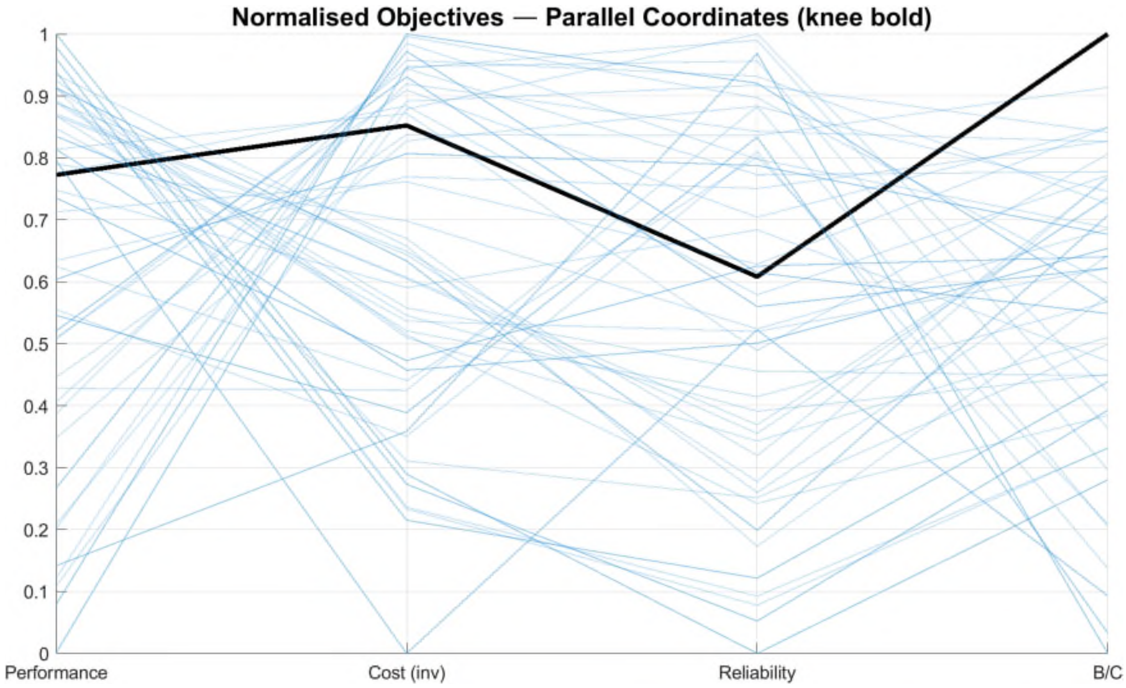

Figure 4-18 Parallel-coordinates plot of normalised objectives; the knee solution is bolded

Figure 4-18, a parallel coordinates plot, enables the comparison of all Pareto-optimal solutions across the four normalised objectives. Each light blue line is a single sensor suite. The bold black line represents the knee solution. It clearly demonstrates that this solution achieves a high level of Performance, Reliability, and Benefit-to-Cost, while requiring only a moderate compromise on Cost (shown as inverted).

The identified knee solution provides a concrete recommendation. It consists of 12 sensors: 5 from the Engine, 2 from the Fuel system, 2 from the EPS, and 3 from the ECS. This suite achieves a normalised performance score of 0.69 at a cost of approximately $36k, with a combined reliability of 145 kh MTBF, as summarised in Table 4-5.

Table 4-5 Composition and objective values for the knee point obtained by multi-objective optimisation

Objective	Knee Value	Description
Diagnostic Performance	≈0.69	Normalised NDCI-based score of the selected suite
Cost	≈US$36k	Sum of sensor purchase and integration costs
Reliability	≈145 kh	Harmonic mean of sensor MTBFs
Sensors per Subsystem	Engine 5, Fuel 2, EPS 2, ECS 3	Composition of the knee suite

From subsystem validation to the platform recommendation (how Table 4-5 is obtained):

Per-subsystem evidence.

Sensors are ranked within each subsystem using the NDCI values; these rankings provide the diagnostic value of candidate sensors.

Platform design space and objectives.

A feasible catalogue of multi-subsystem SENSORS ARTRIBUTIONS is hypothetically generated under integration rules and subsystem quotas. Each suite is scored on the three primary objectives: Performance (↑NDCI-based diagnostic score aggregated over selected sensors), Cost (↓purchase + integration), and Reliability (↑suite MTBF as defined below). A derived view, Benefit-to-Cost (↑a normalised combination of performance and reliability per dollar), is used only for visual triage and does not define Pareto optimality.

Sensor-attribute catalogue (inputs and ranges used for optimisation):

The optimisation uses a curated catalogue of 96 candidate sensors spanning four subsystems. Across the catalogue:

Costs range from \$240–\$12.5k (medians by subsystem: Engine \$1.55k, Fuel \$1.16k, EPS \$0.65k, ECS \$0.65k).

MTBFs ($\mathrm{MTBFsuite}=1/\sum i(1/\mathrm{MTBF}i)$) span 88–350 kh (medians by subsystem: EPS 177 kh, ECS 151 kh, Engine 148 kh, Fuel 119 kh; harmonic-mean MTBFs: EPS 184 kh, ECS 159 kh, Engine 144 kh, Fuel 120 kh).

By sensor family, typical cost/MTBF ranges are:

• Flow: \$4.0–12.5k, 88–150 kh
• Torque: \$2.54–6.48k, 91–152 kh
• Pressure: \$0.88–2.38k, 105–193 kh
• Temperature: \$0.46–1.19k, 127–219 kh
• Thermodynamic property : \$0.81–1.85k, 101–187 kh
• Electrical (voltage/current/power): \$0.24–1.26k, 161–232 kh
• Computed metrics (e.g., thrust, TSFC): \$300, 350 kh

These ranges (detailed list of the sensor attributes is shared with the code bundle [24]) explain the patterns in the Pareto plots: flow/torque channels tend to be costly and less reliable, temperature/electrical channels are cost-efficient with higher MTBF, and Engine reliability is quantitatively close to ECS (medians 148 kh vs 151 kh), so “EPS/ECS most reliable” should be read as small, catalogue-level advantages, not categorical separation.

Pareto filtering and knee selection.
non-dominated suites form the Pareto set (Figs. 16–17). A knee is chosen at the point of highest trade-off efficiency in normalised objective space (largest local curvature / smallest distance to the utopia corner among neighbours). That single design is summarised in Table 4-5 and detailed in Figs. 19–20.

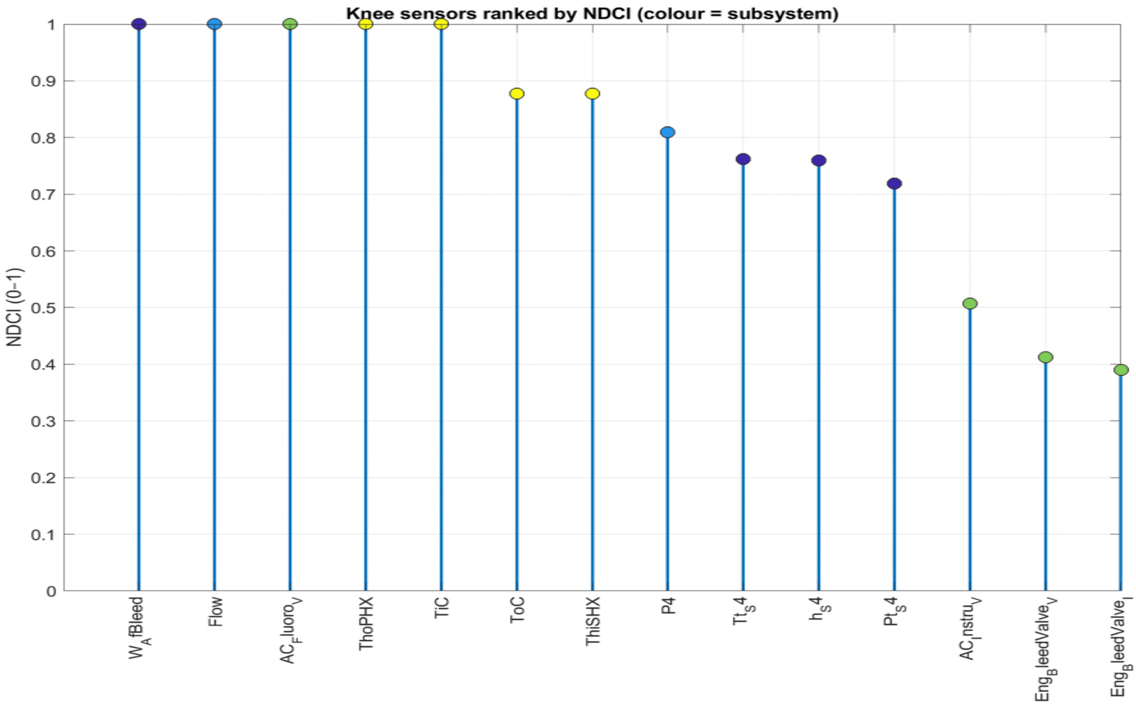

Figure 4-19 The sensors comprising the knee solution, ranked by their NDCI score

The lollipop plot in Figure 4-19 ranks the individual sensors included in the knee solution by their NDCI score. The height of each stem represents the diagnostic value of that sensor. The colours indicate the source subsystem. The mix reflects the catalogue’s economics (with educated guesses): ECS temperature and electrical families provide cost-efficient, reliable coverage, while a limited number of Engine/Fuel flow channels supply targeted separation despite higher cost and lower MTBF.

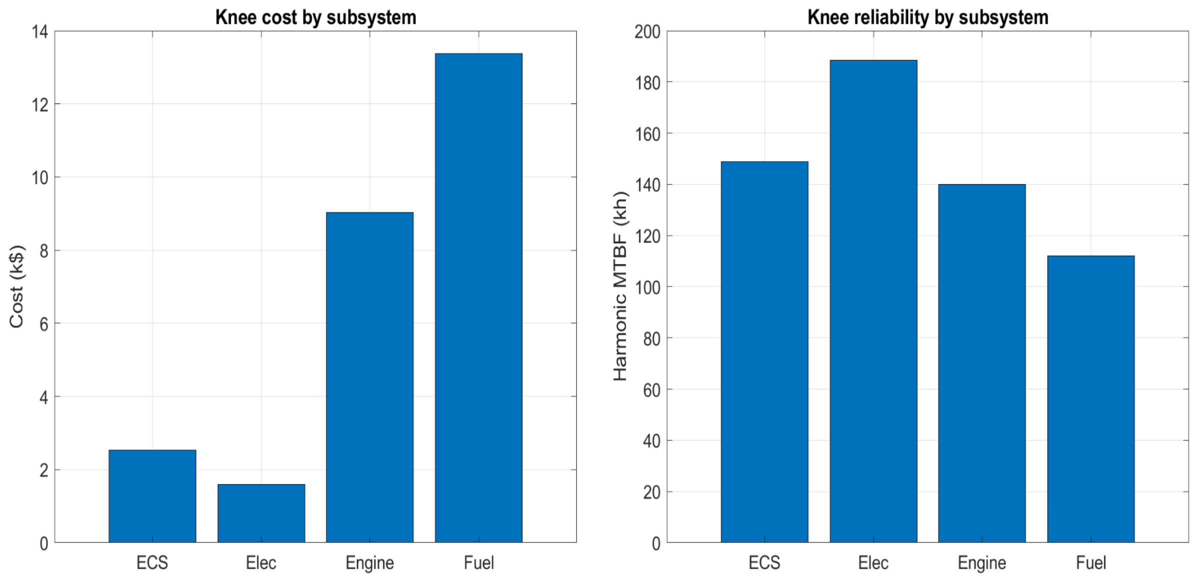

Figure 4-20 Cost and reliability breakdown of the knee suite by subsystem

These bar charts in Figure 4-20 detail the composition of the knee solution. The left panel accumulates cost by subsystem for the twelve selected sensors; higher spending is associated with Engine/Fuel flow families (catalogue costs in the $4–12.5k band), whereas ECS temperature and EPS electrical measurements contribute at lower unit cost.

The right panel reports a descriptive harmonic-mean MTBF for the sensors drawn from each subsystem in the knee. Catalogue medians indicate EPS ≈ 177 kh, ECS ≈ 151 kh, and Engine ≈ 148 kh (Engine is close to ECS), so differences are modest. The suite-level reliability is governed primarily by the lowest-MTBF elements included (typically flow/torque), which explains the overall ≈ 145 kh value in Table 4-5. Please note that the complete sensor-attribute table (including per-sensor cost and MTBF values) is provided with the code bundle [24] to ensure reproducibility without overloading the main text.

4.5 Discussion

The results show that a domain-specific ranking metric, NDCI, yields more informative and compact diagnostic suites than generic relevance-based methods, particularly mRMR, ANOVA, and Information Gain. The advantage is most visible in engine and EPS subsystems, where sensor counts are higher: NDCI's Uniqueness term suppresses redundant channels within a cluster, so the selected set spans complementary loci of the physical process rather than multiple views of the same point.

The knee suite reflects these macro couplings: a small number of high-leverage flow/enthalpy channels (fuel and bleed-air) are combined with temperature chain and electrical measurements that respond downstream. Generic relevance methods (e.g., mRMR/ANOVA) tend to cluster within a single subsystem's highly correlated family (e.g., multiple adjacent temperatures). In contrast, NDCI's Uniqueness term distributes picks along the physical pathway—from cause to propagation—yielding better isolation with fewer sensors.

The Pareto front provides a map rather than a single prescription. A cost-constrained operator can select within the \$32–36 k band, where performance is relatively flat. A certification-focused OEM can move toward points with higher reliability at an additional cost. Maintenance organisations can choose the lowest-cost feasible suites and accept a slight reduction in detection/isolation capability. The 3D view clarifies why designs with similar performance and cost separate in practice; reliability differs and helps avoid choices that raise performance while inadvertently depressing reliability.

The dataset is simulated; validation on rig or fleet data is essential to capture noise, drift, and maintenance effects. Cost and MTBF inputs come from a curated attribute catalogue (vendor figures where available, otherwise documented class averages); suite reliability is computed as a series-equivalent MTBF (reciprocal of summed failure rates), which appropriately penalises the weakest elements but ignores shared-cause failures. Extending the redundancy model beyond linear correlation, injecting prognostic value as an additional objective, and stress-testing robustness to data missingness are natural next steps.

The MOSOF results translate these technical findings into actionable, stakeholder-centric intelligence. The Pareto front is not just a collection of solutions but a strategic decision-making tool. An airline can use the recommended knee point to implement a cost-effective diagnostic upgrade. An Original Equipment Manufacturer (OEM), who may prioritise system reliability for certification purposes, could select a different point on the Pareto front that offers a higher MTBF at the cost of a few additional sensors. An MRO provider might prioritise a suite with the absolute lowest cost, accepting a slight reduction in diagnostic accuracy. By delivering the entire Pareto front, MOSOF empowers each stakeholder to make informed decisions based on their unique constraints and priorities without needing to re-run the entire optimisation analysis.

In sum, the combination of platform-level symptom evidence, properly nested subsystem validation, and multi-objective optimisation explains both the numerical results and their physical meaning: the preferred suites are those that trace the distinguished faults identified with the NDCI score while meeting cost and reliability constraints.

4.6 Conclusions

This paper presents a refined and rigorously validated framework for multi-objective sensor optimisation, built on the foundation of the diagnostically aware NDCI metric. The comprehensive, cross-subsystem evaluation demonstrated that the NDCI-driven approach consistently produces more compact and effective sensor suites for complex diagnostic tasks than conventional methods, such as mRMR. The use of a nested cross-validation protocol ensures that the reported performance metrics are robust and reliable.

By integrating NDCI within the MOSOF, a powerful tool has been created that translates complex engineering trade-offs into a clear set of Pareto-optimal solutions. The application to an airline case study yielded a practical, cost-effective sensor suite that leverages synergistic information from across multiple subsystems. This work represents a significant step towards a more systematic, evidence-based methodology for designing the next generation of IVHM systems, enabling a shift from reactive to predictive maintenance and enhancing the safety and efficiency of complex machinery.

Designing diagnostic sensor suites for integrated aircraft requires balancing performance, cost, and reliability while capturing cross-subsystem fault propagation. The study establishes a complete and reproducible pathway for this task by combining (i) platform-level vectors to expose macro effects, (ii) rigorous, nested per-subsystem validation to obtain unbiased performance estimates of NDCI, and (iii) a MOSOF that delivers stakeholder-ready trade-offs.

Key findings are as follows.

Diagnostic ranking: Across subsystems, the NDCI (which scores sensors by separation power, severity sensitivity, and uniqueness) consistently produces more compact and practical suites than relevance-based baselines. Under the nested protocol, NDCI outperforms mRMR on Engine (balanced accuracy 0.886 vs. 0.690) and ECS (0.677 vs. 0.520), is comparable on EPS (0.518 vs. 0.510), and concedes a small margin on Fuel (0.487 vs. 0.530), where limited sensor diversity constrains uniqueness. These results align with the stepwise inner CV curves, which show that NDCI reaches target accuracy with fewer sensors, indicating improved sample efficiency.

Cross-subsystem evidence: Evaluation on PSV and severity sweeps of fault modes reveals physically plausible macro pathways; for example, Fuel to Engine via combustion signatures, Engine bleed-air to ECS along the heat-exchanger temperature chain, and ECS to EPS via load changes were all visible on the associated PSVs. FM5’s PSV is presented in section 3.1 for illustrative purposes; other PSVs can be obtained from the code bundle provided in the references. NDCI’s uniqueness term disperses selections along these pathways, avoiding clusters of near-duplicate measurements. This explains the higher isolation capability with compact suites.

Stakeholder trade-offs and the Pareto front: Multi-objective optimisation over Performance↑ - Cost↓ - Reliability↑ yields a feasible Pareto set with clear regimes: a low-cost, flat-performance floor; a moderate-slope region; and a diminishing-returns region near 0.69–0.71 performance. The knee solution—identified by maximum local curvature in normalised objective space—comprises 12 sensors (Engine 5, Fuel 2, EPS 2, ECS 3) delivering ≈ 0.69 performance at ≈ US$36k with ≈ 145 kh suite MTBF. The 3D Pareto view clarifies why designs with similar performance and cost can differ materially (reliability is the key factor that separates them), supporting informed choices for airlines, OEMs, or MROs.

Reliability treatment: Suite-level reliability is computed as a series-equivalent MTBF (the reciprocal of the summed failure rates), which appropriately penalises the weakest elements. Catalogue medians show only modest differences among subsystems (e.g., the Engine is close to the ECS), so overall suite reliability is driven by the specific sensor families selected (e.g., flow/torque tend to be costlier and less reliable than temperature/electrical).

The integrated NDCI–MOSOF workflow translates raw multi-sensor data into transparent, defensible evidence for IVHM design. It avoids optimistic bias through proper nesting, exposes the physical meaning of selections via component-wise ranking visuals, and delivers a portfolio of optimal suites rather than a single prescription, enabling stakeholder-specific decisions without re-running the complete analysis.

Limitations and future work. Findings are derived from a high-fidelity virtual aircraft and curated cost/MTBF catalogues. Validation on rig and in-service fleet data is required to account for noise, drift, maintenance actions, and shared-cause failures. Future extensions include (i) richer redundancy models beyond linear correlation, (ii) incorporation of prognostic value (e.g., RUL) as an additional objective, (iii) robustness to missing data and varying mission profiles, and (iv) uncertainty-aware optimisation of cost and reliability inputs.

Overall conclusion. A diagnostically aware ranking coupled with principled, multi-objective optimisation provides a traceable, platform-level route to sensor suites that capture cross-subsystem effects while respecting economic and reliability constraints. The approach is readily transferable to other complex assets. It supports model-based safety and certification processes by delivering reproducible, decision-grade artefacts, from stepwise validation to Pareto fronts and a clearly justified knee recommendation.

4.7 References

1. Suslu, B.; Ali, F.; Jennions, I.K. Understanding the Role of Sensor Optimisation in Complex Systems. Sensors 2023, 23, 7819.
2. Suslu, B.; Ali, F.; Jennions, I.K. Normalised Diagnostic Contribution Index (NDCI) Integration to Multi Objective Sensor Optimisation Framework (MOSOF)—An Environmental Control System Case. Sensors 2024, 24, 2661.
3. Jennions, I.K. Integrated Vehicle Health Management: Perspectives on an Emerging Field; SAE International: Warrendale, PA, USA, 2011.
4. Ezhilarasu, C.M.; Skaf, Z.; Jennions, I.K. A Generalised Methodology for an Effective Diagnosis of Aircraft Systems with Interacting Subsystems. IEEE Access 2021, 9, 59616–59633.
5. Jia, L.; Ezhilarasu, C.M.; Jennions, I.K. Cross-Condition Fault Diagnosis of an Aircraft Environmental Control System (ECS) by Transfer Learning. Appl. Sci. 2023, 13, 12686.
6. Vachtsevanos, G.; Lewis, F. L.; Roemer, M.; Hess, A.; Wu, B. Intelligent Fault Diagnosis and Prognosis for Engineering Systems; John Wiley & Sons: Hoboken, NJ, USA, 2006.
7. Zhang, W.; Wang, J. A framework of airplane Integrated health management. In Proceedings of the 2016 IEEE Second International Conference on Multimedia Big Data (BigMM), Taipei, Taiwan, 20–22 April 2016; pp. 363–366.
8. Peng, H.; Long, F.; Ding, C. Feature selection based on mutual information: Criteria of max-dependency, max-relevance, and min-redundancy. IEEE Trans. Pattern Anal. Mach. Intell. 2005, 27, 1226–1238.
9. Lei, Y.; Li, N.; Guo, L.; Li, N.; Yan, T.; Lin, J. Machinery health prognostics: A systematic review from data acquisition to RUL prediction. Mech. Syst. Signal Process. 2018, 104, 799–834.

10. Civera, M.; Surace, C. A multi-objective genetic algorithm for robust optimal sensor placement. Comput.-Aided Civ. Infrastruct. Eng. 2021, 36, 1419–1436.
11. Mezura-Montes, E.; Coello Coello, C.A. Multiobjective Evolutionary Algorithms in Aeronautical and Aerospace Engineering. IEEE Trans. Evol. Comput. 2014, 18, 493-514.
12. Deb, K.; Pratap, A.; Agarwal, S.; Meyarivan, T. A fast and elitist multiobjective genetic algorithm: NSGA-II. IEEE Trans. Evol. Comput. 2002, 6, 182–197.
13. Zitzler, E.; Thiele, L. Multiobjective evolutionary algorithms: a comparative case study and the strength Pareto approach. IEEE Trans. Evol. Comput. 1999, 3, 257–271.
14. SAE International. ARP4754B: Guidelines for Development of Civil Aircraft and Systems; SAE International: Warrendale, PA, USA, 2023.
15. Federal Aviation Administration. AC 25.1309-1B, System Design and Analysis; FAA: Washington, DC, USA, 1996.
16. What the True Cost of Aircraft Components Reveals About Industrial Supply Chains , https://impactograph.com/what-the-true-cost-of-aircraft-components-reveals-about-industrial-supply-chains/
17. Structural Health Monitoring Cost Estimation of a Piezosensorized ..., https://pmc.ncbi.nlm.nih.gov/articles/PMC8915022/
18. Integrated cost/weight optimization of aircraft structures Markus Kaufmann · Dan Zenkert · Per Wennhage
19. Optimal Sensor Selection for Health Monitoring Systems - NASA Technical Reports Server (NTRS), https://ntrs.nasa.gov/api/citations/20050237898/downloads/20050237898.pdf
20. Multi-objective optimization using genetic algorithms: A tutorial ..., https://pure.psu.edu/en/publications/multi-objective-optimization-using-genetic-algorithms-a-tutorial

21. Chapman, J. W., Lavelle, T. M., May, R. D., Litt, J. S., & Guo T. T-MATS: Toolbox for the Modeling and Analysis of Thermodynamic Systems. 2014. Available at: https://doi.org/10.2514/6.2014-3932
22. Jennions, I.K.; Ali, F. An Assessment of Heat-Exchanger Degradation in a B737-800 Environmental Control System (ECS) Using the System Environment for SESAC (SESAC). J. Therm. Sci. Eng. Appl. 2021, 13, 051010.
23. Andrews, D.K.; Jennions, I.K.; Cordelia, E.A. A Platform View of Aircraft Data Simulation and Analysis, 2025. Available at doi:10.20944/preprints202503.1945.v1
24. B. Suslu, Code bundle - https://github.com/ssl8/NDCI-with-MOSOF

5 Conclusion, Contributions, and Directions for Future Research

This thesis addressed the critical and persistent challenge of designing effective health management systems for complex, safety-critical aerospace platforms. The research was motivated by the significant operational and economic pressures within the commercial aviation industry, where the effectiveness of Maintenance, Repair, and Overhaul (MRO) activities directly impacts safety, reliability, and financial viability. It was posited that the foundation of any advanced Aircraft Health Management System (AHMS) is its sensing infrastructure, and that the systematic optimisation of this infrastructure represents a key-enabling step towards more predictive and efficient maintenance paradigms.

The overarching aim of this research was to establish and validate a novel, systematic framework for the multi-objective optimisation of sensor networks to enhance the health management of modern civil aircraft. This chapter concludes the thesis by summarising how this aim was achieved through the defined research objectives. It then explicitly lists the principal contributions to knowledge that have resulted from this work and, finally, identifies promising directions for future research that can build upon the foundations established herein.

5.1 Conclusion and Fulfilment of Objectives

The thesis set out to establish and validate a principled pathway for multi-objective sensor optimisation that measurably improves diagnostic detection and isolation for complex, safety-critical aircraft. The central premise was that health management capability is limited by the information quality of the sensing layer, and that rigorous optimisation must therefore couple diagnostic value, system-level trade studies, and stakeholder constraints. The work delivered this premise in three stages: (i) a structured synthesis of the state of the art and a gap analysis (Chapter 2); (ii) the design of a Multi-Objective Sensor Optimisation Framework (MOSOF) with an NDCI-centred performance metric and a multi-objective genetic algorithm solver (Chapters 3–4); and (iii) empirical validation on high-fidelity subsystem and platform-level studies with repeated, nested cross-validation for unbiased estimation (Chapter 4).

Fulfilment of objectives.

Objective 1 - Theoretical foundation and taxonomy. Chapter 2 systematised sensor optimisation across selection, placement, data processing, and operation; formalised objective/cost functions for each; and identified a methodological gap: a lack of diagnostically aware, system-of-systems optimisation under explicit multi-objective trade-offs.

Objective 2 - MOSOF design. A modular framework was formulated to integrate performance, cost, and reliability objectives, and to expose Pareto-efficient sensor suites under real constraints.

Objective 3 - NDCI development. The Normalised Diagnostic Contribution Index (NDCI) was defined to quantify a sensor's separation power, severity sensitivity, and informational uniqueness, thereby aligning selection with fault detection and isolation needs rather than purely statistical relevance.

Objective 4 - Subsystem application. An ECS case (Chapter 3) demonstrated that NDCI-guided suites achieve similar or higher diagnostic capability with fewer sensors than relevance-based baseline mRMR.

Objective 5 - Cross-subsystem validation. A platform-level study (Engine, Fuel, EPS, ECS) demonstrated that NDCI-MOSOF leverages fault-propagation pathways to assemble compact, cross-subsystem suites, which were validated through repeated nested cross-validation to mitigate optimistic bias.

Objective 6 - Evaluation and limitations. Strengths, constraints, and transferability were analysed, with recommendations for real-rig/fleet validation, robustness to missingness, uncertainty treatment, and certification-aligned evidence generation.

Answers to the research questions.

RQ1 (multi-objective optimisation for AHMS). The thesis presents a repeatable MOO process in which MOGA evolutionary search explores non-dominated trade-offs among diagnostic performance, cost, and reliability; knee-point selection yields defendable choices for different stakeholders.

RQ2 (quantifying diagnostic contribution). The NDCI aggregates diagnostically meaningful attributes (separability, severity response, uniqueness) into a transparent index that prioritises sensors with system-level impact rather than merely high correlation to labels.

RQ3 (benefit across interconnected subsystems). When applied across Engine, Fuel, EPS, and ECS, the NDCI-MOSOF pipeline identifies suites that detect and isolate both local and propagated faults using fewer channels than relevance-based baselines, while honouring cost/reliability limits.

Quantitative summary and decision-grade artefacts. Under consistent learners and repeated *nested* validation, NDCI-guided suites proved to be compact and practical. Across subsystem-level evaluations, NDCI generally outperformed mRMR while producing compact sensor suites, with the Fuel subsystem being the main exception because of limited sensing diversity. At platform level, MOSOF produced a defendable 12-sensor knee solution balancing diagnostic performance, acquisition cost, and reliability. These results, stepwise accuracy curves (inner CV), aggregate confusion matrices (outer tests), and multi-objective fronts, provide traceable evidence for design and certification dossiers.

Methodological strengths. Three features underpin the scientific rigour: (i) diagnostic alignment via NDCI (physics- and task-aware ranking); (ii) unbiased performance estimation via repeated *nested* cross-validation; and (iii) explicit trade studies via multi-objective Pareto analysis and knee-point justification.

Scope and limitations. Findings were derived from a high-fidelity virtual aircraft and curated catalogue attributes; rig/fleet data and operational noise (drift, maintenance actions, shared-cause failures) remain to be incorporated. Reliability was treated as a series-equivalent MTBF; extensions to common-cause and phase-of-flight dependent reliability are warranted. Nonetheless, the consistency of advantages across subsystems, together with the compactness of suites and the decision-grade Pareto evidence, supports the central claim: diagnostically aligned, multi-objective sensor optimisation delivers smaller, stronger, and more defensible sensing solutions for aircraft health management.

5.2 Contributions

This research yields several original contributions to knowledge in the fields of sensor optimisation and Integrated Vehicle Health Management (IVHM). The principal contributions, spanning theoretical, methodological, and empirical domains, are summarised:

1. Theoretical and methodological.

- A diagnostically aware ranking index (NDCI) that unifies separability, severity, sensitivity, and uniqueness into a single, normalised measure of sensor value for fault detection and isolation.
- A general MOSOF formulation that couples NDCI-based performance with cost and reliability objectives, solved with evolutionary MOO (MOGA) and equipped with knee-point selection for stakeholder-specific decision-making.
- A validated evaluation protocol through four aircraft subsystems analysed using repeated nested cross-validation to eliminate optimistic bias in feature/sensor selection studies.
- A coverage-preserving selection procedure (scenario-wise minimal sets and informed unions) to respect rare-but-critical faults while avoiding oversensing.

2. Empirical and practical.

- Subsystem-level evidence (ECS, Engine, Fuel, EPS) that NDCI-guided suites attain higher balanced accuracy with fewer sensors than relevance-based baselines, except where sensor diversity is intrinsically low (Fuel).
- Platform-level trade studies producing a defendable Pareto set and a knee suite (12 sensors) with ≈ 0.69 performance at $\approx$ US$36k and ≈ 145 kh MTBF, evidence directly usable by OEMs, airlines, and MROs.
- Reproducible, decision-grade artefacts (rank decompositions, confusion matrices, stepwise curves, Pareto/parallel-coordinates plots) that expose

why specific sensors are chosen and how they trade against cost and reliability.

3. Positioning within the literature.

- The thesis operationalises community recommendations on diagnosability-driven selection and system-of-systems reasoning in IVHM, while clarifying when relevance-redundancy heuristics such as mRMR underserve diagnostic needs.

5.3 Recommendations for Future Work

The results obtained across Chapters 3-4 demonstrate that prioritising network-level diagnostic worth through the Normalised Diagnostic Contribution Index (NDCI) and embedding that metric within the Multi-Objective Sensor Optimisation Framework (MOSOF) produce compact, defensible sensor suites with measurable utility at the aircraft level. The knee solution identified for platform deployment shows that diagnostic gain can be achieved without unsustainable growth in cost or complexity. Figures 4-16 to 4-20 further illustrate a transparent trade space that stakeholders can interrogate during the design review. Building on these findings, the following research directions are proposed to consolidate scientific validity, expand operational relevance, and strengthen certifiability. They are expressed as a single, integrated agenda to preserve the style and structure of this thesis section.

1- Platform-grade validation should be extended from simulation-driven campaigns to heterogeneous operational datasets. The repeated nested cross-validation protocol already mitigates optimistic bias; however, external validity can be further enhanced by curating multi-fleet, multi-route corpora that span climatic regimes, maintenance practices, and avionics baselines. Blocked evaluations (by aircraft tail, by day of operation, by environmental regime) can quantify generalisation under domain shift, while hold-out campaigns from unseen fleets can measure transferability. Particular attention should be given to the Fuel subsystem, where mRMR marginally outperformed NDCI under limited diversity, to determine whether additional or alternative sensing modalities (e.g., differential pressure across filters, temperature upstream/downstream of the FOHE) or refined symptom construction can recover separability when failure signatures are weak or sparse. Where ethically and practically feasible, Hardware-in-the-Loop and seeded-fault ground tests should be introduced to anchor digital-twin results in controlled physical responses.

2- The NDCI should be extended from a primarily static, score-based index to a dynamic, uncertainty-aware diagnostic value measure. Three refinements are recommended: (i) time-resolved NDCI that aggregates separation, severity sensitivity, and uniqueness over operational phases and transients, so that a sensor's contribution is recognised when and where it matters (take-off, climb, descent, hot-day ground turnarounds); (ii) probabilistic NDCI, in which aleatoric and epistemic uncertainties in features, labels, and models are propagated to deliver confidence-weighted contribution scores; and (iii) groupwise NDCI that quantifies complementarity and conditional redundancy within sensor tuples, ensuring that apparent single-sensor utility is not double-counted once correlated peers are selected. Together, these extensions will enable MOSOF to reward sensors that are diagnostically decisive only under specific operating contexts, penalise fragile contributions, and favour robust, synergistic suites.

3- Diagnosability-driven design should be integrated more closely with the platform symptom vector (PSV) abstraction. While PSV has proven beneficial to encode cross-subsystem manifestations, explicit isolation guarantees at the aircraft level remain an open target. Future work should (i) compute subsystem and platform diagnosability matrices that encode which fault pairs can be separated by a candidate suite, (ii) embed these matrices directly into the NDCI or the multi-objective fitness as constraints (e.g., minimum fault isolation ability for all safety-critical pairs), and (iii) model propagation across mechanical, thermal, and pneumatic interfaces using causal graphs so that MOSOF allocates sensors where ambiguity persists at boundaries (e.g., engine-ECS bleed interactions). Such a programme would turn the current “best-effort” isolation into a quantified, requirement-traceable capability, directly supporting certification arguments.

4- Robust and risk-aware optimisation should be explored to complement the current Pareto-efficient search. Airline and MRO stakeholders tolerate cost and reliability variations differently from OEMs; moreover, sensor performance and failure rates vary with age, maintenance practices, and environmental conditions. Chance-constrained, distributionally robust, or worst-case formulations can ensure that target detection and isolation levels are met under plausible shifts in noise, missingness, or sensor drift. Likewise, the cost objective should evolve from acquisition-centric figures to lifecycle-accurate models that incorporate installation labour, calibration, spares pooling, dispatch penalties, and maintenance-credit opportunities. Performing post-optimal sensitivity analysis on these variables will reveal whether the knee point is locally stable or whether small budget or reliability perturbations cause the suite composition to change, which is critical for sustainment planning.

5- The placement dimension should be brought into the same optimisation loop as selection and operation. This thesis treated selection as primary, with placement discussed at the review level; however, placement choices influence both the signal-to-noise ratio and survivability (e.g., routing vulnerability, thermal stress). Future MOSOF variants should treat location as a decision variable, subject to installation constraints, access for maintenance, and wiring/power budgets. Where analytic gradients are unavailable, multi-fidelity surrogates, which combine CFD/thermodynamic models for ECS ducting, heat-transfer estimates for engine bays, and structural transmissibility maps for vibration, can provide adequate placement guidance at a tractable cost. The objective set would then explicitly trade detection/isolation gain against added mass, harness length, and installation risk, further grounding the front in airframer reality.

6- Real-time deployment constraints should shape both the data pipeline and the learning stack used for evaluation. On-aircraft compute and bandwidth limits, downlink policies, and on-condition maintenance triggers necessitate lightweight, certifiable algorithms and deterministic execution. Subsequent studies should therefore (i) profile candidate classifiers and fusion schemes for latency and memory on representative avionics hardware, (ii) employ streaming-first feature engineering with bounded state, and (iii) implement graceful degradation under missing or late data. The diagnostic policies derived from NDCI-MOSOF must be re-verified under these constraints to ensure that fleet-level benefits survive integration.
7- Explainability and decision transparency should be expanded from sensor ranking to case-level diagnostic reasoning. NDCI already delivers interpretable per-sensor values; however, line maintenance and engineering need mechanisms that trace a given alert or isolation outcome back to a minimal, human-readable set of signals and signatures. Future work should develop counterfactual and perturbation-based justifications ("which readings, if removed or altered within tolerance, would flip the decision?") and map these onto maintenance manual tasks and troubleshooting trees. The result would be a continuous thread from Pareto-front design decisions to on-the-ramp actions, increasing trust and shortening AOG durations.
8- Standards alignment and certification pathways should be explicitly targeted so that optimisation outputs are admissible as evidence. Harmonisation with open information models and condition-monitoring data exchange standards will reduce integration friction and facilitate multi-stakeholder adoption. In parallel, safety cases should formalise detection/false alarm targets for hazardous failure conditions, demonstrate the required availability and integrity of the sensing chain, and document configuration control processes for NDCI and MOSOF artefacts. The framework is well placed to supply the quantitative

backbone of such cases; formalising the claims, arguments, and evidence will accelerate acceptance by airworthiness authorities.

9- Cross-platform generalisation should be investigated through structured transfer. The framework should be tested on additional aircraft types, propulsion architectures, and environmental control architectures to examine invariances in ranking and suite composition. When full retraining is impractical, domain adaptation and parameter-sharing strategies, guided by physical similarity rather than purely statistical fit, can maintain diagnostic value with limited labelling effort. This line of work is directly relevant for airline groups operating mixed fleets and for MROs seeking toolchain commonality.

10- Resilience, security, and sustainability should be treated as first-class optimisation concerns. Sensor cyber-hardening (including spoofing and tamper detection), the secure provenance of data used for NDCI estimation, and the environmental footprint of sensor procurement and disposal (including rare-earth content and recyclability) are increasingly material to operators and regulators. Future MOSOF variants should incorporate constraints or objectives that reflect these realities and explicitly quantify trade-offs on the same fronts as cost, reliability, and diagnostic benefit.

Finally, the research programme should culminate in open, reproducible assets and practitioner tooling. Releasing de-identified datasets, reference implementations of NDCI-MOSOF (including nested CV protocols), and curated case studies would enable independent verification and benchmarking. Internally, a decision-support interface should be developed that exposes the Pareto set, allows users to interrogate sensitivity to weights and constraints, and surfaces the “why” behind each candidate suite in a manner aligned with engineering and maintenance practice. By coupling transparent methods, rigorous validation, and actionable tooling, the framework can evolve from a promising research demonstrator to an operational capability that consistently delivers predictive maintenance value across fleets and years in service.