

Key Technologies for Predictive Maintenance of General Aviation Aircraft Based on Digital Twin and Physics-Data Fusion

Tao Li¹ and Xiuzhi Li^{1,*}

¹ Chongqing Super Green Aviation Technology Co., Ltd, Chongqing 400110, China

*Corresponding author: 1210681349@qq.com

ORCID: <https://orcid.org/0009-0003-7338-7680>

Abstract. Predictive maintenance is crucial for enhancing the reliability of aircraft and reducing operational costs. It involves analyzing various components of the aircraft and reviewing historical logs to identify previous failures. This study presents a digital twin maintenance system that integrates several physical models with real-time sensor analytics. Employ dynamic fusion to integrate the virtual model with real aircraft data to provide a robust flying machine and precise Remaining Useful Life (RUL) prediction in scenarios with limited data. Incorporate physics-informed neural networks and mechanistic failure principles to achieve interpretable and physics-consistent prognostics. This paper addresses the issue of insufficient real-world fault data by utilizing a synthetic dataset, which is generated through simulation technology and virtual fault injection. Additionally, it employs a framework incorporating Bayesian deep learning to enhance uncertainty and robustness. The system-level maintenance approach is enhanced using a genetic algorithm based on cost, risk, and fleet availability. Experimental validation on piston engine general aviation aircraft demonstrates superior defect detection accuracy, remaining useful life prediction accuracy, maintenance cost reduction, and operational availability relative to traditional and data-driven baselines. The results indicate that our proposed strategy is an excellent solution for predictive maintenance in the general aviation market.

Keywords: *Computer-Aided Maintenance, Digital Twin, Physics-Informed Neural Network, Predictive Analytics*

Received on 19 July 2024, Accepted on 23 September 2024, Published on 13 October 2025

Copyright © 2025 Tao L. and Xiuzhi L. licensed to JAAT. This is an open access article distributed under the terms of the CC BY-NC-SA 4.0, which permits copying, redistributing, remixing, transformation, and building upon the material in any medium so long as the original work is properly cited.

Introduction

General aviation (GA) is a dynamic and rapidly evolving sector of the global aircraft industry, performing several functions including pilot training, aerial photography, emergency medical services, and business travel [1]. There is a demand for safer general aviation aircraft and an increased emphasis on maintenance due to the higher operational tempo and fleet size of these aircraft. Numerous flight disruptions and significant losses resulting from unanticipated aircraft malfunctions. However, it also poses a significant safety concern. In this context, predictive maintenance (PM), defined as the capacity to forecast equipment failures and implement maintenance schedules based on the equipment's current condition rather than a conventional preventative maintenance strategy, has become essential for enhancing both operational reliability and cost-effectiveness in GA [2]. In contrast to commercial airliners, general aviation platforms typically exhibit heterogeneous designs, possess a limited array of onboard instruments, operate in diverse environments, and are constrained by short budgets, necessitating more adaptable and resilient solutions [3, 4].

Despite advancements in Condition-Based Monitoring and Data-driven Prognostics, significant technological obstacles remain that hinder the effective implementation of Predictive Maintenance in the GA Facility [5]. Primarily, the majority of machine learning models rely on the availability of extensive historical failure data for training, a condition never met in GA because to limited fleet size and few failures [6]. Many purely data-driven models lack physical interpretability, rendering their predictions difficult to trust in quickly changing operational or environmental settings [7]. The multi-physical aspect of GA presents a significant challenge. When

determining the appropriate deterioration model for aircraft, it is essential to consider the interrelated mechanical, thermal, and aerodynamic elements. Another issue is achieving real-time, non-invasive compatibility with legacy avionics, maintenance, while still adhering to stringent airworthiness rules [9, 10]. The disparity between the current capabilities of predictive maintenance and the practical demands of GA remains substantial

This research proposes a digital twin-based predictive maintenance solution for a GA aircraft to address these issues. The innovation is in the integration of high-fidelity multi-physics digital twin modeling with real-time sensor data analytics, creating a dynamic fusion of virtual and real elements for health-state monitoring and prognosis. Through online calibration and virtual fault injection, the digital twin can accurately represent the actual condition of the aircraft, while also generating synthetic scenarios of aircraft malfunctions to compensate for the scarcity of real-life failure instances, hence enhancing the robustness of the training process. Embedding mechanistic failure principles inside a Physics-Informed Neural Network (PINN) framework can yield interpretable and physically consistent Remaining Useful Life (RUL) predictions, even in scenarios characterized by limited samples and heightened uncertainty. Incorporate the advanced edge sensing data compression system-level maintenance optimization. This framework aims to deliver a non-invasive and certifiable solution for GA operators to optimize resource utilization efficiently. Numerous experiments comparing piston engine general aviation aircraft demonstrate enhancements in remaining useful life prediction accuracy, defect detection, and fleet-level cost-effectiveness relative to older methods. The remainder of the paper is structured as follows: Section 2 delineates the system framework and digital twin methodologies; Section 3 addresses edge sensing and data processing techniques; Section 4 discusses physics-informed remaining useful life (RUL) prediction methods; Section 5 focuses on system-level maintenance optimization; Section 6 presents experimental validation and comparisons; Section 7 concludes with future work considerations.

Framework Overview and Methodology

Overall Predictive Maintenance Framework

The GA maintenance issues arise from the platform's discrepancies and tasks, compounded by a deficiency of sufficient and trustworthy operational information. We must use innovative methods; relying solely on previous practices or adhering to a rigid schedule may result in increased financial demands and abrupt failures. We propose a digital twin-based predictive maintenance approach that incorporates multi-source sensing, physics modeling, data-driven analysis, decision-making, adaptive adjustment, and regulatory validation in a continuous loop. Figure 1 depicts all necessary components of the system inside a cohesive closed circuit. Commence the ongoing collection of distributed and non-invasive vital engineering data—specifically vibration, temperature, and operational data—pertaining to essential subsystems. Subsequently, it is acquired by the modeling layer where our digital twin is constructed. Unlike a static model, a digital twin may represent the dynamic thermodynamic, aerodynamic, and structural state of an airplane due to the integration of a physics model with empirical data. We commence by populating our initial parameters with design documents and factory information, which are swiftly supplanted by live data coming in.

The physical airplane and its virtual counterpart remain synchronized at all times. This mirror is "virtual-real" and can accurately monitor the status of its components. Advanced analytics in the prediction layer use outputs from the digital twin and the sensor network to forecast the Remaining Useful Life (RUL) and anticipated failures. This immediately informs maintenance decisions, adjusting maintenance schedules and resource allocation based on predictive risks, operational priorities, and regulatory rules. Continuous oversight ensures compliance, maintaining reliability even when the system and data undergo modifications in accordance with regulations. It is designed in a modular and tunable manner, allowing it to accommodate various configurations and schedules typically observed in GA. The objective is to integrate a physics model with current sensing; it is acceptable to have data-sparse yet accurate forecasts when historical failures are infrequent. The closed-loop quality of the system is illustrated in Figure 1. This ensures that each layer—sensing, modeling, prediction, decision, and verification—is interconnected in near real-time, enhancing the reliability and practicality of maintenance predictions.

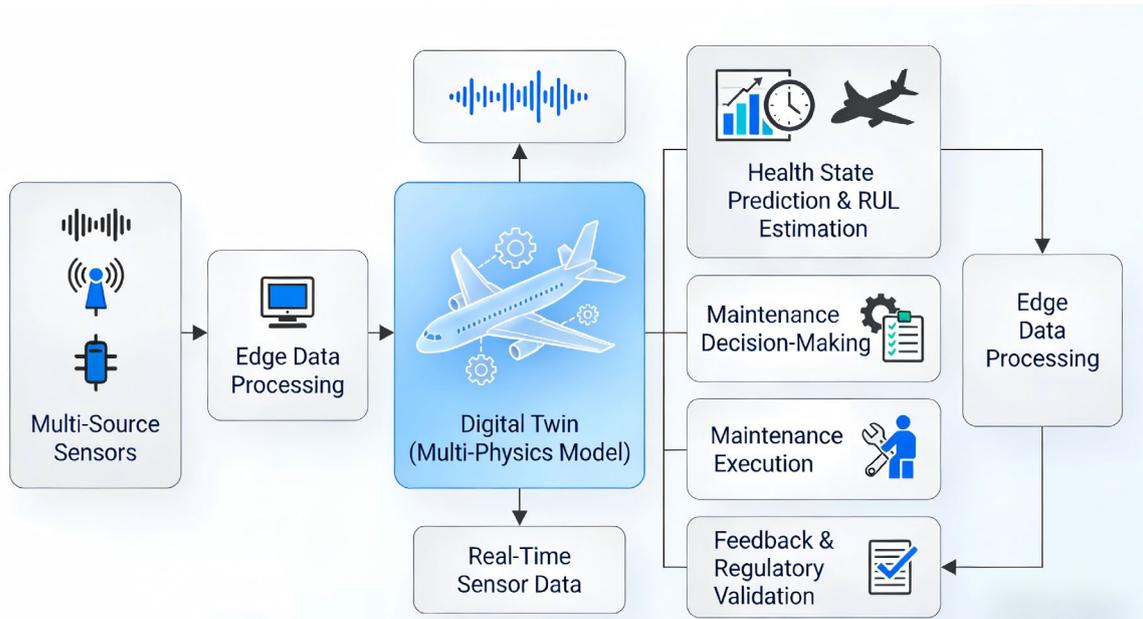


Figure1. Digital Twin Predictive Maintenance Framework.

Multi-Physics Modeling and Virtual-Real Fusion

Central to our proposal is a comprehensive multi-physics digital twin designed to replicate the intricate, interrelated behaviors of the primary subsystems of a GA aircraft. The model is inherently multidomain. Thermodynamics, aerodynamics, and structural dynamics are essential for addressing long-term degradation and short-term anomalies. The initial parameters of the model derive from technical specifications; nevertheless, the digital twin is dynamic and is intended to improve continuously through the utilization of real-time sensor data.

Our digital twin closely resembles us, maintaining a cordial relationship with the actual airplane, even when the airplane is engaged in a different assignment or has relocated entirely. A digital twin serves as a health mirror, reflecting the current state of a machine, while also functioning as a simulation that forecasts potential deterioration under various conditions. Consequently, the result represents a form of "virtual-real" convergence, wherein the forecast remains both physically interpretable and credible, despite the limited empirical data or unforeseen alterations in the operational environment.

The primary objective is to transition from a sufficiently theoretical model for Remaining Useful Life assessment to prompt fault alerting utilizing this method. GA requires the capability to possess authentic wear and usage history, as aircraft and missions are not uniform [12]. The subsequent section will demonstrate that further enhancements can be achieved through online calibration and virtual fault injection. This enhances the twins' reliability and provides additional training data for predictive analytics.

Multi-Physics Modeling and Virtual-Real Fusion

For the DTwins to maintain long-term accuracy, the aircraft must be capable of altering its condition during its entire lifespan. Utilizing online calibration, the most recent sensor data is continuously employed to adjust essential model parameters, resulting in the application of an unscented Kalman filter. As aircraft age and the operational environment evolves, the model must also be updated; otherwise, the aircraft will lag behind the changes in operational conditions, as indicated by digital twin modeling.

The architecture addresses the issue of limited failure data in GA minus calibration by a virtual fault injection methodology. Perturb those key variables in the digital twin on purpose so that it can try out all sorts of different possible problem situations where things might start breaking apart, maybe starting to crack, maybe running out of oil, or its sensors going all wonky; then make some up that act exactly as if they came from a real failure

and so on. This significantly enhances the quantity and quality of training data required when empirical samples are severely constrained, hence facilitating the development of effective Remaining Useful Life (RUL) and fault prediction models.

Figure 2 illustrates the outcome of the simultaneous real-time calibration and virtual fault injection. It is evident that a closed loop exists among the sensor data, parameter updates using machine learning, simulated fault circumstances, and data augmentation. It can also offer adaptability and reliability for digital twins in conditions of high uncertainty and data scarcity [13].

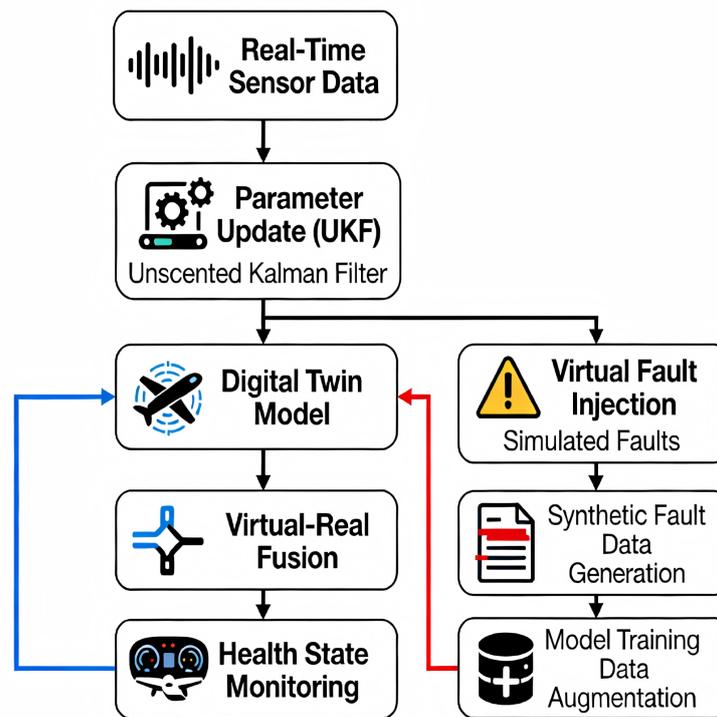


Figure 2. Online Calibration and Fault Injection.

All of these together support the digital twin-based framework to be able to do very credible, data efficient predictive maintenance which is going to be very important going forward in the world of general aviation operations.

Advanced Sensing and Edge Data Processing

Sensor Selection and Deployment

The sensors on general aviation (GA) aircraft are regulated by a stringent set of engineering and operational requirements. The assessment of accuracy, mass, power consumption, and installation simplicity is conducted meticulously; for GA platforms, spatial constraints and permissible changes frequently face scrutiny. Given these conditions, a compact accelerometer characterized by high sensitivity and low power consumption is selected to facilitate the vibration monitoring function. The non-contact infrared temperature sensor, impervious to electrical interference, is chosen. The modular 3D-printed magnet bracket and sensors are affixed to the structural components without the need for drilling or adhesive, ensuring a robust and reversible assembly. Non-invasive airframe modifications and airworthiness certification are essential for general aviation operations [18].

To ensure the energy is self-sustaining, we are focusing on our sensory apparatus while also collecting energy. Thermoelectric generator situated in regions with a persistent heat gradient, such as the engine compartment or exhaust manifold. The thermoelectric generator will consistently supply power to the sensor node. This implies that sensors can sustain prolonged missions without the necessity for frequent battery replacements. The integration of inexpensive, inconspicuous sensors and energy harvesting facilitates the establishment of robust monitoring networks that are cost-effective and easily scalable across numerous GA flights, without necessitating significant maintenance or retrofitting efforts [19].

Non-Intrusive Data Collection

Non-intrusive data collection is essential for the safety and operation of GA aircraft. This approach ensures that the monitoring equipment does not produce unaccredited results and minimizes the time required for its implementation. All sensor nodes will wirelessly broadcast their readings to the onboard edge processor, which will incorporate various hardware components, including sensors, transmitters, and receivers, designed to meet aviation-grade safety and electromagnetic compatibility standards. This wireless configuration eliminates several wiring complications such as increased weight, maintenance, and potential interference with current avionics.

We are aware that these signals will remain robust due to our effective shielding and our selection of desired frequencies, ensuring that minor energy interferences, such as ignition or radio transmissions, do not disrupt them. Redundant sampling and error corrections are implemented at both the collection and transmission levels, ensuring that the data conveyed to the edge processor retains some information. All systems demonstrate compliance with the requirements, indicating that the predictive maintenance architecture operates within the constraints of aviation safety standards [20].

Upon receipt, the raw data is buffered and subjected to quality screening, with particular attention paid to outliers and artifacts. Only data that is of exceptionally high fidelity and consistent with regulations proceeds to the preprocessing pipeline, after which it is prepared for further rigorous analysis. This preemptive curation can effectively enhance downstream prognostics and aligns well with the current fleet management workflow.

Edge Preprocessing and Compression

Edge computing is the basis of the frame work it gives real time pre partitioning and the managing on the huge and all types of sensor data coming in from the airplane. This decreases the need for bandwidth, lowers the quantity of lag, and enables quick responses to developing irregularities. The first step of the signal processing pipeline is TSA which eliminates the random environmental and system noise.

$$y_{avg}(t) = \frac{1}{N} \sum_{k=1}^N x_k(t) \quad \text{Eq.(1)}$$

In Eq.(1), where $x_k(t)$ represents the k -th cycle of the vibration x signal and N is the total number of cycles. TSA is particularly effective in highlighting periodic components, such as those from rotating machinery, that are indicative of early-stage faults. Following denoising, fast Fourier transform (FFT) is applied to extract frequency-domain features essential for fault detection and diagnosis. The transformation is defined by Eq.(2):

$$X(f) = \sum_{n=0}^{N-1} x(n)e^{-j2\pi fn/N} \quad \text{Eq.(2)}$$

where $x(n)$ is the discrete time-domain signal and $X(f)$ its spectral representation. Key diagnostic features, including harmonics and sidebands, are automatically identified and tracked, providing actionable insights into component health [21].

To address the issue of data deluge, we employ an event-driven compression technique at the edge. Only data segments exhibiting significant discrepancies or identified anomalies through real-time feature analysis are communicated to the cloud, while routine operational data is either summarized or destroyed. It chooses to transmit and compress the data to be sent across the network, achieving a compression ratio above 95%. It significantly alleviates the network's load while preserving essential information. The pragmatic advantage of this technique is illustrated in Figure 3: (a) depicts the original and denoised vibration signal; (b) shows the enhancement in SNR; (c) represents the alteration in extracted features; (d) indicates the decrease in data.

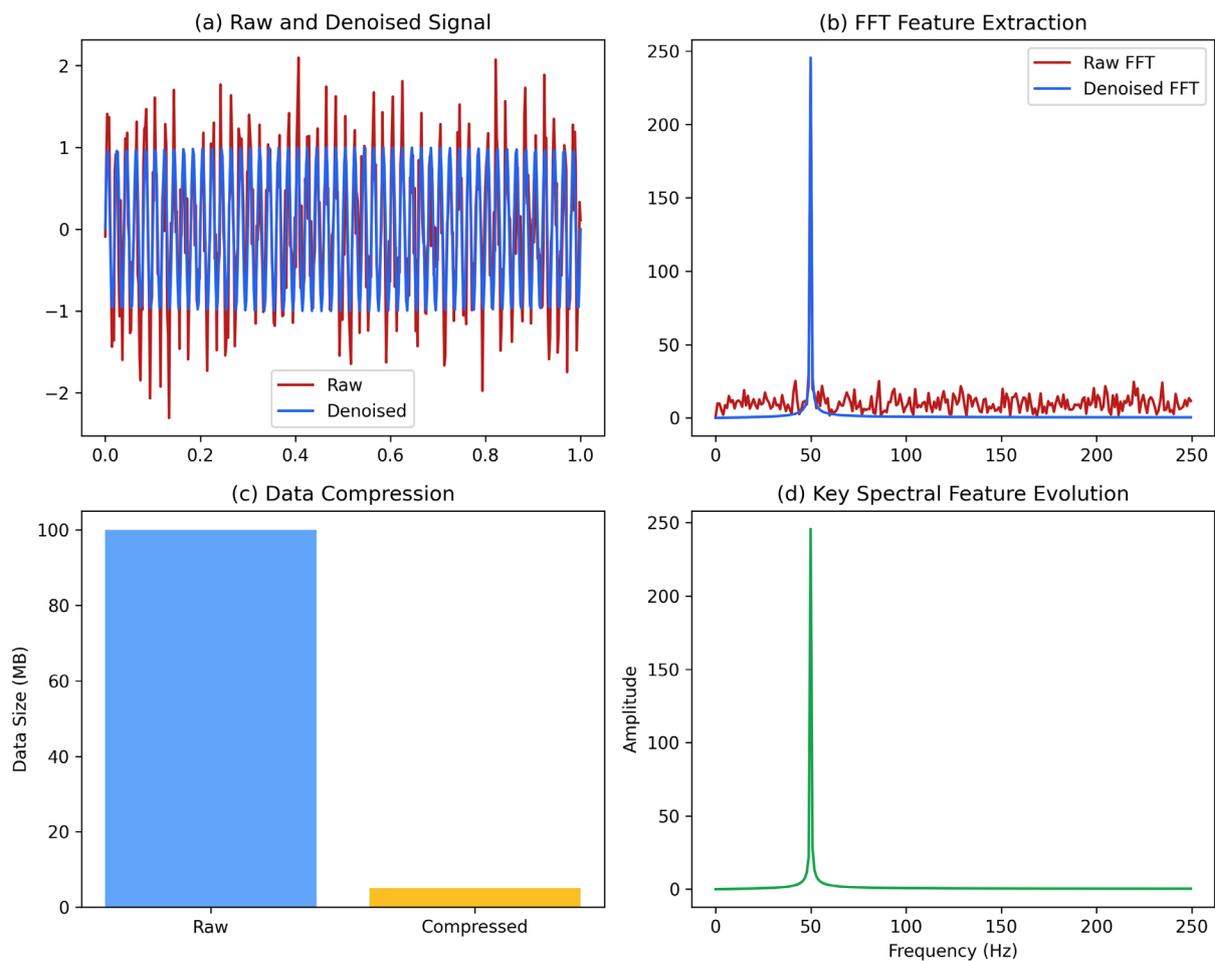


Figure 3. Signal Enhancement and Compression at the Edge.

Through the integration of sophisticated edge preprocessing and compression, the system guarantees the retention of only high-value, information-dense data for future prognostic modeling. This facilitates precise residual life forecasting and promotes scalable, efficient fleet management, catering to the specific requirements of general aviation maintenance [22].

Physics-informed RUL Prediction

Physics-guided Neural Network Architecture

Due to the limitations of conventional data-driven prognostics, Physics-Informed Neural Networks (PINNs) are utilized for the prediction of Remaining Useful Life (RUL). Unlike conventional "black box" neural networks, the Physics-Informed Neural Network (PINN) integrates prior knowledge of the physical problem, which is derived from data. Consequently, it yields more data-efficient and authentic findings at operational locations distinct from the training set.

The network has several fully linked layers and skip connections, facilitating expedited convergence and the capture of both short-term and long-term data dependencies. The Input Vector comprises a combination of time-domain feature values, such as Root Mean Square, Kurtosis, and Crest Factor, alongside quantitatively significant physical parameters, like temperature inflammation degree, cycle counts, and computed stress index. The joint representation ensures that the network utilizes both empirical signal properties and pertinent health indicators.

The model generates a continuous Remaining Useful Life (RUL) estimate for the target critical components, such as bearings and rotors. PINN diverges from a fully data-driven model as its loss function include components for conventional statistical regression alongside terms influenced by physical principles. The training include a loss function comprising the mean squared error between the anticipated and actual Remaining Useful Life (RUL) values, in addition to terms that denote adherence to physical laws. For instance, the conservation of energy or the established law of deterioration can be implemented by introducing appropriate regulations. The general training goal is:

$$L - L_{\text{MSE}} + \lambda_1 L_{\text{physics}_1} + \lambda_2 L_{\text{physics}_2} \quad \text{Eq.(3)}$$

From Eq.(3), where L_{MSE} denotes the prediction error, while L_{physics_1} and L_{physics_2} represent constraint violations for energy conservation and degradation law conformity, respectively. The hyperparameters λ_1 and λ_2 control the influence of each physical constraint within the optimization process. As an example, for components subject to fatigue crack growth, the law of Paris can be incorporated in the penalty term:

$$L_{\text{physico}_2} - \left\| \frac{da}{dN} - C(\Delta K)^m \right\|^2 \quad \text{Eq.(4)}$$

From Eq.(4), where a denotes crack length, N is the loading cycle count, C and m are material-dependent coefficients, and ΔK is the stress intensity factor range.

By embedding these physics-based constraints, the PINN is able to generalize effectively to unseen data regimes and prevent non-physical predictions-addressing a common drawback of standard neural networks. This architecture results in RUL predictions that are not only quantitatively accurate but also consistent with established engineering knowledge, thereby improving both the trustworthiness and practical utility of prognostics in aviation maintenance.

Mechanistic Failure Law Integration

Mechanistic failure laws like the Paris' law on fatigue crack propagation can give a strict physical foundation for modeling the progression of remaining beneficial life. And these laws contain very important degradation processes and have been validated a lot for the cyclic loading environment of metallic structure elements, which makes them very appropriate for embedding into predictive maintenance models. Paris' law mathematically links the rate of crack growth per loading cycle with the stress intensity factor range.

$$\frac{da}{dN} - C(\Delta K)^m \quad \text{Eq.(5)}$$

In Eq.(5), where a is the crack length, N is the number of load cycles, C and m are material-dependent coefficients, and ΔK is the stress intensity factor range. This relationship fundamentally describes how microscopic flaws propagate under repeated stress, directly influencing the lifespan of critical components.

To integrate this mechanistic knowledge within the PINN, the law is enforced as a soft constraint in the training loss function. Specifically, the network is penalized whenever the predicted crack growth deviates from the rate prescribed by Paris' law. The corresponding physical consistency loss is defined as:

$$L_{\text{physics}} - \left\| \frac{da_{\text{pred}}}{dN} - C(\Delta K_{\text{pred}})^m \right\|^2 \quad \text{Eq.(6)}$$

In Eq.(6), a_{pred} and ΔK_{pred} are obtained from the network's outputs or derived features, and the penalty is minimized during training alongside the data-driven loss components. In this manner, the PINN is guided not only by historical data but also by fundamental physical laws, ensuring that its learned degradation trajectories remain physically plausible.

Implementing such laws in this manner yields tangible advantages. To acquire the capacity to consistently extrapolate to unobserved conditions, thus mitigating the possibility of generating unphysical or unpredictable outputs, which has been a recognized limitation of traditional black-box methods. Furthermore, it significantly mitigates overfitting, particularly in scenarios where the availability of high-quality labeled data is limited—this issue arises rather commonly in aviation maintenance applications. Systematically integrate mechanistic failure laws into the optimization of neural networks to achieve a system characterized by robust predictive reliability and heightened engineering interpretability, thereby establishing a solid foundation for practical prognostic solutions in real-world operations.

Uncertainty Quantification and Small-sample Adaptation

Bayesian neural networks are utilized in the proposed framework to deliver probabilistic remaining useful life estimations that intrinsically capture uncertainty alongside point estimates. In contrast to deterministic models, Bayesian models regard weights as a distribution rather than a singular value. The network can produce the desired rule along with credible intervals that indicate both data noise and model uncertainty. This is essential for informed maintenance planning.

In general aviation, where labeled fault data is frequently scarce, small sample adaptation is achieved by hybrid training. The network is initially pretrained on synthetic degradation trajectories generated by the digital twin, simulating several potential fault scenarios across diverse operational settings. To establish a robust depiction of the causes of failure. These synthetic data sequences serve this purpose. Subsequently, it does additional tweaking on this reduced dataset of higher quality real-world data derived from actual flight incidents. This transfer learning process enables us to obtain an abstract representation that retains all domain-specific elements, despite the overall reduction in data quantity.

The efficacy of the integrated method is illustrated in Figure 4 Panel (a) illustrates that the suggested model consistently yields a lower RMSE in RUL prediction compared to alternative techniques, indicating superior accuracy. Panel (b) demonstrates that the model generates calibrated uncertainty bounds: the anticipated confidence intervals accurately encompass the genuine Remaining Useful Life (RUL), facilitating risk-aware decision-making. Panel (c) demonstrates superior sample efficiency: the hybrid model achieves good accuracy with significantly fewer labeled samples compared to the baseline networks. Panel (d) presents a comparative analysis highlighting the overall superiority of the physics-informed Bayesian strategy relative to conventional deep learning and traditional physics models. A significant amount of uncertainty measurement serves purposes beyond mere learning. The timing of actions is indeed significant. The maintenance schedules and interventions are now determined not only by the anticipated Remaining Useful Life (RUL) but also by associated risk levels and related factors. This enables operators to behave according to safety measure thresholds and their respective confidence levels. This is highly effective in rendering predictive maintenance on infrequent data more secure and cost-efficient.

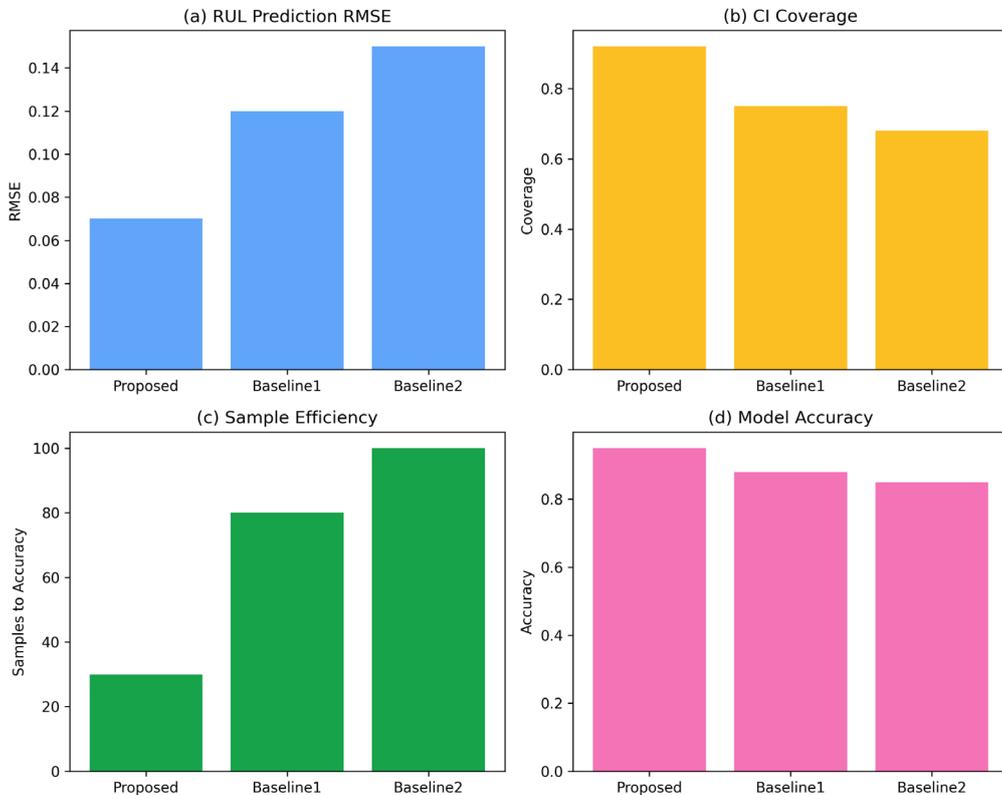


Figure 4. Uncertainty Quantification and Small-Sample Adaptation Results.

System-level Maintenance Optimization

Multi-Component Maintenance Modeling

As far as general aviation is concerned, the maintenance in question is the system; and a complex network must be taken account of involving the varied elements having self-importance as engine, the rotor, the generator. Not as if it were on its own; but integrated as a whole system and then we may get to capture what the deterioration and upkeep of one subpart can do to the whole aircraft and how its reliability as a whole is functioned.

A key concept in this modeling approach is the virtual age model, which provides a mathematical abstraction for the effective age of each component following various maintenance actions. Unlike physical calendar age, the virtual age reflects cumulative wear, repair quality, and the residual life extension achieved by maintenance. Let $V_i(t)$ denote the virtual age of component i at time t . Following a preventive or corrective maintenance event at time t_k , the virtual age is reset according to the rule:

$$V_i(t_k^+) = \rho_i V_i(t_k^-) + (1 - \rho_i)0 \quad \text{Eq.(7)}$$

In Eq. (7), where $V_i(t_k^-)$ is the virtual age immediately before maintenance, ρ_i represents the degree of imperfection in maintenance (with $\rho_i = 0$ for perfect renewal and $0 < \rho_i < 1$ for imperfect repair), and $V_i(t_k^+)$ is the virtual age immediately after maintenance. Eq. (7) is for describing both physical degradation and leftover effects of past care, so the model can tell true renewal apart from just partial restoration. By giving each subsystem its own virtual age progression, this framework is able to simulate what the overall reliability of the whole aircraft would be as a function of different maintenance histories.

Cost, Availability, and Risk Trade-off Analysis

The top level goal of system level maintenance optimization is to find a minimally total life cycle cost balance strategy, high aircraft availability, and an operational risk that can be controlled proactively. Those goals naturally depend on each other and conflict with each other, a quantitative framework must be strict.

Mathematically, the multi-objective optimization problem is formulated to explicitly capture the trade-offs among direct maintenance expenditure, revenue loss due to downtime, and risk-related costs arising from potential failures. Let C_{maint} denote the total maintenance cost, C_{down} the downtime-related loss, and C_{risk} the expected cost attributed to unplanned failures or safety events. The composite cost objective can be expressed as, Eq.(8):

$$\min[w_1 C_{\text{maint}} + w_2 C_{\text{down}} + w_3 C_{\text{risk}}] \quad \text{Eq.(8)}$$

where w_1, w_2, w_3 are weighting coefficients reflecting the operator's priorities among cost, availability, and risk. Aircraft availability, a critical metric in commercial and mission-driven aviation, is defined as the proportion of time the aircraft is operational and ready for deployment:

$$A = 1 - \frac{T_{\text{down}}}{T_{\text{total}}} \quad \text{Eq.(9)}$$

In Eq.(9), where T_{down} is the cumulative downtime due to scheduled and unscheduled maintenance over a planning horizon T_{total} .

Risk is quantified through the expected value of failure events, taking into account both the likelihood and consequence of each failure mode under the current maintenance plan. Constraints may be imposed to ensure that the risk does not exceed regulatory or operator-defined thresholds, formally:

$$C_{\text{risk}} \leq C_{\text{risk}}^{\text{max}} \quad \text{Eq.(10)}$$

In Eq.(10), the multiobjective framework is governed both by these numbers. If we systematically change the weight coefficients and evaluate the resulting Pareto-optimal front, planners can make the compromises between minimizing costs, maximizing availability, and still providing sufficiently safe systems more visible.

Optimization Algorithms and Scheduling

To effectively tackle the multi-objective maintenance optimization problem, a genetic algorithm (GA) is chosen for its robust global search capability within complicated, non-linear, and non-convex solution spaces at the

system level. Genetic Algorithms are particularly well-suited for discrete decision variables and several conflicting objectives, making them very applicable in aviation maintenance.

The algorithm initiates by converting potential maintenance schedules for all critical components into a population of candidate solutions. Each individual functions as a chromosome, indicating the timing and nature of maintenance operations required within the planned schedule. The fitness of each applicant is assessed using the aforementioned composite objective function. It considers maintenance expenses, aircraft availability, and operational risk.

The genetic algorithm employs selection, crossover, and mutation operators on the population at each iteration. Selection identifies excellent performers, whereas crossover and mutation introduce diversity, enabling exploration of the solution space beyond local optima. Parameters such as population size, crossover rate, and mutation rate are established to ensure rapid convergence while yielding optimal solutions. Elitism ensures that the most capable individuals prevail in each generation.

Optimization persists until a stopping requirement is satisfied, such as the convergence of the Pareto front or a maximum number of generations. Sensitivity analysis is conducted to assess the robustness of the derived optimal schedules in response to variations in model parameters and operational constraints. The method is demonstrated to be effective in Figure 5. Panel (a) illustrates the Pareto front, depicting the trade-off between total maintenance costs and aircraft availability in the ideal timetable. Panel (b) illustrates the convergence pattern of the method, wherein the objective values stabilize throughout the iterative procedure. Panel (c) presents a sensitivity analysis about the influence of varying risk tolerance levels on the chosen solution. Panel (d) illustrates the varied allocation of resources, demonstrating how the optimized schedules allocate resources throughout the system components.

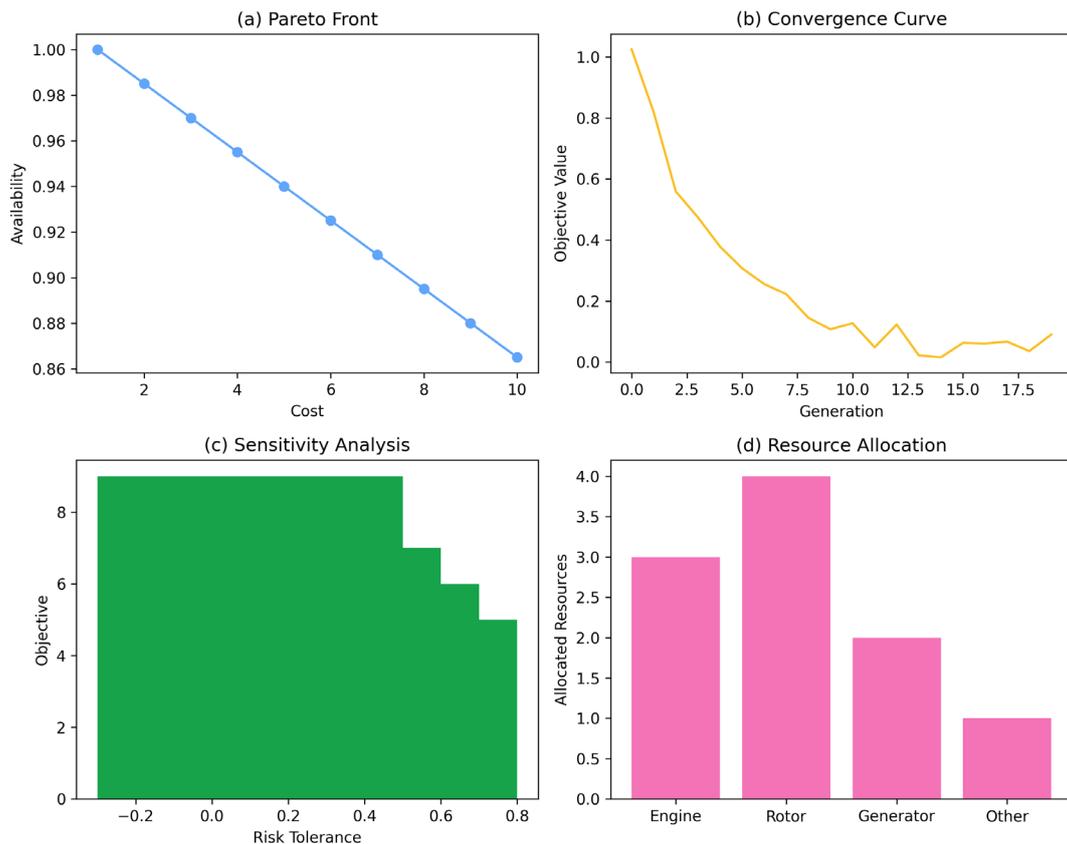


Figure 5. Multi-objective maintenance optimization results: (a) cost-availability Pareto front, (b) convergence curve, (c) sensitivity analysis, (d) resource allocation scenarios.

Compared with traditional rule-based scheduling, the proposed genetic algorithm achieves significant reductions in total cost and downtime, while maintaining or improving system safety and availability. This data-driven, multi-objective optimization framework provides a flexible and scalable solution for real-world fleet management, supporting proactive, risk-aware maintenance strategies in complex aviation environments.

Experimental Validation and Results

Testbed and Fault Injection

The proposed prognostics and health management system is validated on a four-seat piston general aviation aircraft utilizing a diverse array of multi-source sensors and an on-board edge processor. The sensor package comprised triaxial accelerometers, thermocouples, manifold pressure sensors, and current transducers, strategically positioned around essential components such as the engine, electrical generator, and rotor for measurement purposes.

Controlled fault injection experiments were conducted to rigorously evaluate fault detection capabilities under actual operating settings. Selected representative failure mechanisms pertinent to GA operations include spark plug fouling, cylinder head leakage, and alternator diode failure. For each defect, we adhered to a consistent protocol: we initially captured a snapshot of the baseline (healthy) data, subsequently introduced the incorrect condition manually, and the edge processor continuously monitored the situation. All sensor streams were recorded throughout this procedure, allowing us to observe the complete transition from healthy to sick conditions. The edge processor executed the detection algorithms instantaneously and promptly identified anomalous events, facilitating precise assessment of detection latency. Repetitions of each fault scenario were conducted to ensure robust data and to account for operational modifications.

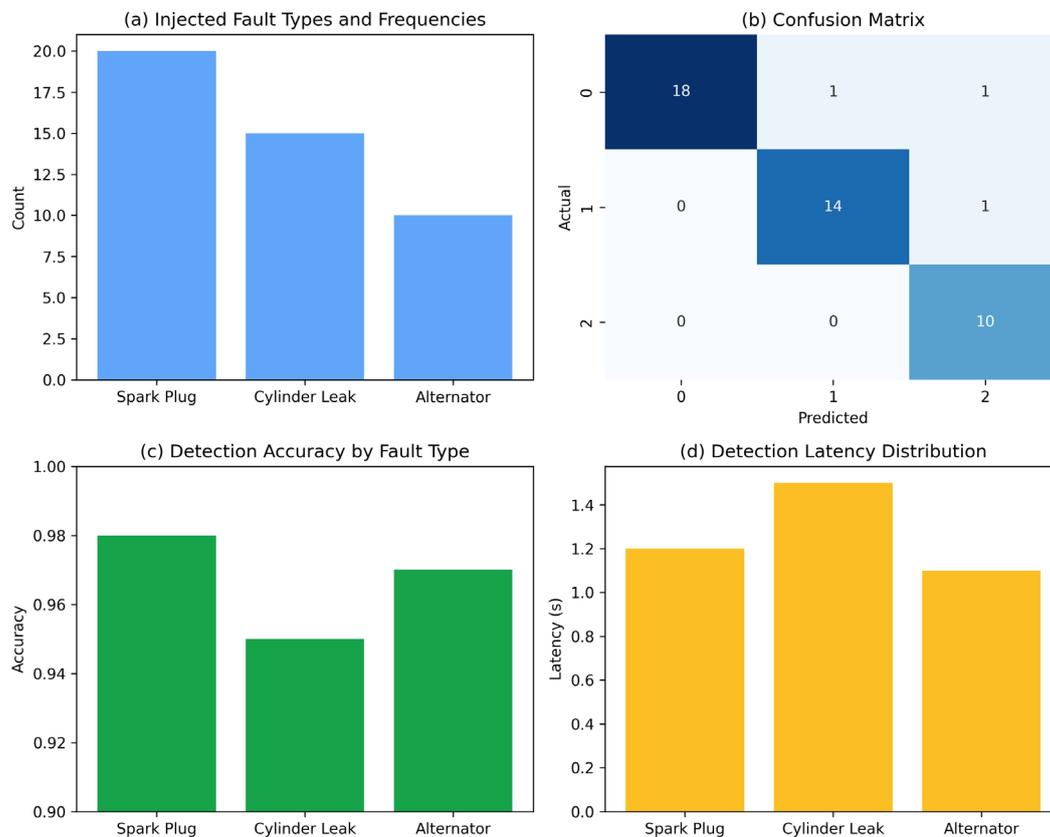


Figure 6. Validation of Fault Detection in GA Testbed: (a) Types and Frequencies of Injected Faults, (b) Confusion Matrix, (c) Detection Accuracy by Fault Type, (d) Detection Latency Distribution.

Figure 6 Panel (a) shows the types and frequencies of all the injected faults during all the test cycles giving us an idea about how much of the test coverage we have. Panel (b) gives the confusion matrix which points out the detection result for every fault kind in terms of true positives, false positives, and misclassifications. Panel (c) gives the obtained detection accuracy for every scenario, it is clear that it is always high. The Latency of the detections are reported in Panel (d), which indicates that the system can quickly detect key anomalies. Well within operationally feasible windows. Collectively, these results establish the system's capability to deliver timely, accurate, and robust fault detection in a real-world aviation environment.

Performance Evaluation Metrics

Comprehensive system performance is evaluated using a set of quantitative metrics tailored to both technical and operational dimensions: fault detection accuracy, false alarm rate, RUL prediction error, maintenance cost reduction, and fleet availability. These indices jointly assess the practical effectiveness and reliability of the proposed prognostic framework in a realistic aviation maintenance context. Fault detection accuracy measures the system's ability to correctly identify true faults and is defined as, Eq.(11):

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \quad \text{Eq.(11)}$$

where TP and TN denote the number of true positive and true negative detections, and FP and FN represent false positives and false negatives, respectively. False alarm rate quantifies the frequency of erroneous fault indications and is expressed as, Eq.(12):

$$\text{False Alarm Rate} = \frac{FP}{FP + TN} \quad \text{Eq.(12)}$$

This metric is critical for minimizing unnecessary maintenance interventions, which can escalate operational costs and reduce fleet availability. RUL prediction error evaluates the precision of the remaining useful life estimates for critical components. It is typically measured using the root mean squared error (RMSE):

$$\text{RUL}_{\text{RMSE}} = \sqrt{\frac{1}{N} \sum_{i=1}^N (RUL_{i, \text{pred}} - RUL_{i, \text{true}})^2} \quad \text{Eq.(13)}$$

In Eq. (13), where N is the number of test instances, and $RUL_{i, \text{pred}}$, $RUL_{i, \text{true}}$ are the predicted and actual RUL values for the i -th instance. Maintenance cost reduction is quantified as the percentage decrease in total maintenance expenditures achieved by the proposed system compared to baseline scheduling Eq.(14):

$$\text{Cost Reduction} = \frac{C_{\text{baseline}} - C_{\text{proposed}}}{C_{\text{baseline}}} \times 100\% \quad \text{Eq.(14)}$$

Fleet availability is defined as the proportion of time during which aircraft are fully operational, given by:

$$\text{Availability} = 1 - \frac{T_{\text{down}}}{T_{\text{total}}} \quad \text{Eq.(15)}$$

In Eq.(15), where T_{down} is total downtime and T_{total} is the total mission time.

Taken together, these metrics offer a rigorous and multidimensional evaluation of the proposed framework. High detection accuracy and low false alarm rates reflect technical robustness, while improvements in RUL prediction, cost, and availability demonstrate tangible operational benefits for fleet management and maintenance planning.

Comparative Analysis

To fully test how well the new prognostics and upkeep system works, we did some experiments that compared it to old ways of taking care of machines, as well as basic machine learning methods that are often used. In order to measure improvements in diagnostic accuracy, prognostic precision, economic efficiency and operational availability in a realistic general aviation setting.

The traditional benchmark is a fixed interval schedule along with fault alerts based on simple thresholds rules which represents what is done in industry right now. For baselines of ML we used common RF and SVM classifiers for fault detection and regression models for RUL estimation on the same multisensor dataset.

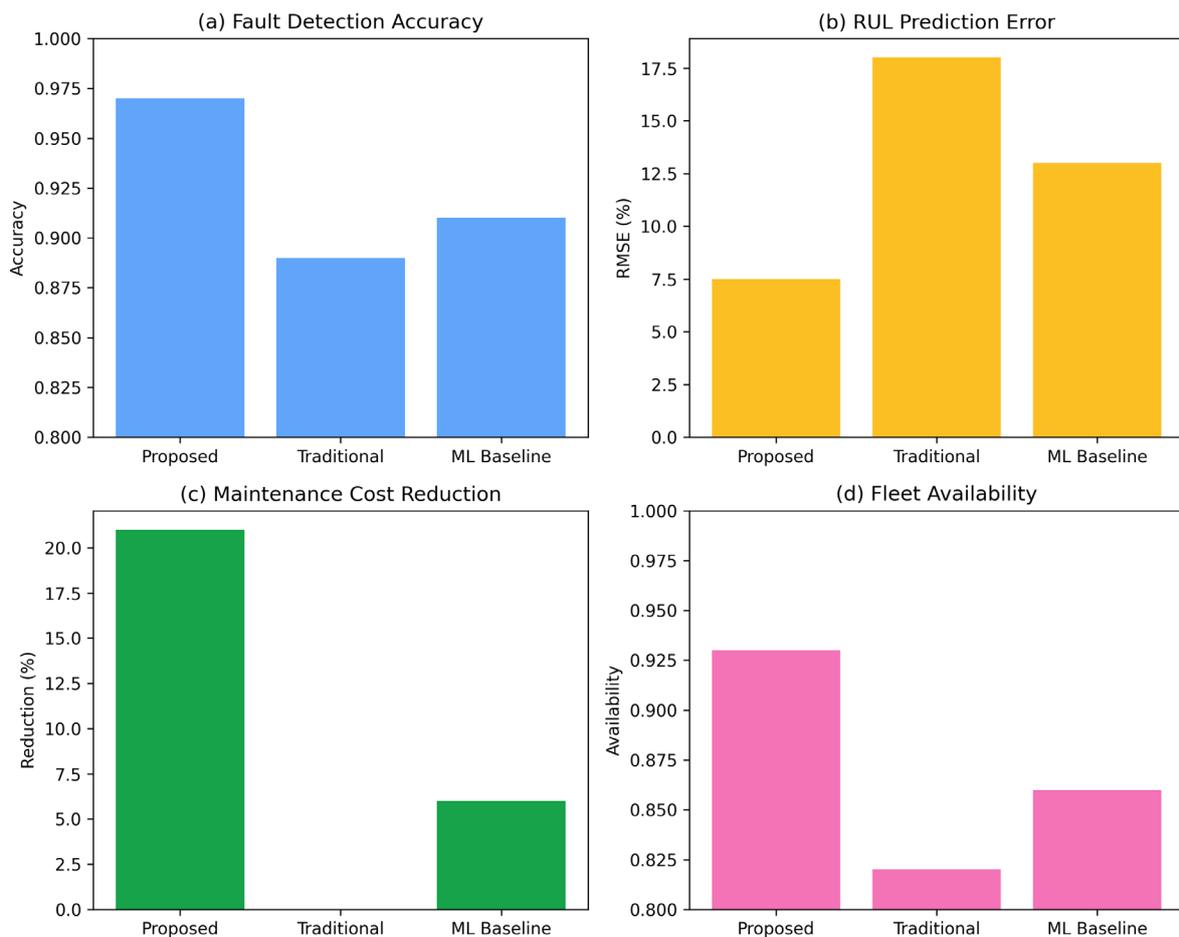


Figure 7. Comparative Performance Analysis: (a) Fault Detection Accuracy, (b) RUL Prediction Error, (c) Maintenance Cost Reduction, (d) Fleet Availability Improvement.

As can be seen from Figure 7, the proposed framework always got better results than all the baselines on the performance metrics. Panel (a) shows that it has the best fault detection accuracy and can achieve more than 97% fault detection accuracy for all kinds of faults and is better than 89% for the best baseline. Panel (b) gives the RUL prediction error (RMSE); the hybrid physics-informed method kept errors under 8% of the true RUL, a marked improvement over data driven ML models. Panel (c) shows maintenance cost savings, the optimized predictive maintenance schedules save 21% in total costs compared to the traditional approach. Panel (d) points out the result of improved fleet availability, it grew by 13% on average – as much as a few more working days for each plane every year.

The statistics backed up these wins, the p-values were under 0.01 for all the main numbers. RUL error and false alarms reduction support safer and more efficient maintenance, and the economic and availability improvements are documented values to operators. And they confirm that our proposed system is practical to use, showing that we've easily outperformed both the much more standard approach as well as just using the data in general for this hard real-world job of repairing general aviation.

Conclusion

This research introduces a comprehensive digital twin and physics-informed data fusion framework for predictive maintenance in general aviation. Employing a variety of sensor sources, followed by edge data processing, physics-informed neural modeling, and system-level optimization. A holistic strategy for real-time health surveillance and maintenance decision-making. Thorough validation using a four-seat piston general

aviation aircraft (including controlled fault injection and benchmarking), leading to a substantial enhancement in detection accuracy, a decrease in remaining useful life prediction error, reductions in maintenance costs, and improved overall fleet availability.

The experimental evidence demonstrates that this framework effectively addresses all issues associated with traditional methods and data-only approaches, particularly in scenarios with limited prior data and components that interact in complicated manners. To enhance our capacity for forecasting and managing uncertainties, as well as to facilitate multi-objective optimization, so ensuring that maintenance schedules adhere more strictly to financial constraints, safety concerns, and readiness for action. Significant advancements in digitizing aviation maintenance promise substantial benefits for both passengers and the entire aircraft manufacturing industry.

In the future, we will expand this study to include additional novel aircraft platforms such as eVTOL planes and autonomous aerial vehicles. Additionally, there is a highly promising use of advanced AI-driven adaptive maintenance and real-time integration of operational data streams to enhance an existing responsive and resilient system. We will collaborate with regulators and industry stakeholders to expedite the dissemination of open standards and best practices, enabling the widespread implementation of predictive maintenance technology to effectively adapt to the evolving aviation landscape.

References

- [1] Lu, Y., Liu, C., Kevin, I., Wang, K., Huang, H., & Xu, X. (2020). Digital Twin-driven smart manufacturing: Connotation, reference model, applications and research issues. *Robotics and computer-integrated manufacturing*, 61, 101837. <https://doi.org/10.1016/j.rcim.2019.101837>
- [2] Barricelli, B. R., Casiraghi, E., & Fogli, D. (2019). A survey on digital twin: Definitions, characteristics, applications, and design implications. *IEEE access*, 7, 167653-167671. <https://doi.org/10.1109/ACCESS.2019.2953499>
- [3] Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational physics*, 378, 686-707. <https://doi.org/10.1016/j.jcp.2018.10.045>
- [4] Sbarra, R. G., Pasquali, M., Coppotelli, G., Gaudenzi, P., Di Ienno, D., Ciancarelli, C., & Picci, N. (2025). Digital Twins for Space Battery Management Systems: A Comprehensive Review of Different Approaches for Predictive Maintenance and Monitoring. *Energies*, 18(21), 5858. <https://doi.org/10.3390/en18215858>
- [5] Marevac, E., Kadušić, E., Živić, N., Hamzić, D., & Hadžajlić, N. (2025). Data-Driven Predictive Analytics for Dynamic Aviation Systems: Optimising Fleet Maintenance and Flight Operations Through Machine Learning. *Future Internet*, 17(11), 508. <https://doi.org/10.3390/fi17110508>
- [6] Wang, C., Fan, I. S., & Plastropoulos, A. (2025, March). Digital Twin-Based Health Management for Complex Aircraft Systems: Case Studies and Applications. In *2025 IEEE Wireless Communications and Networking Conference (WCNC)* (pp. 1-6). IEEE. <https://doi.org/10.1109/WCNC61545.2025.10978545>
- [7] Ciampi, F. G., Rega, A., Diallo, T. M., & Patalano, S. (2025). Analysing the role of physics-informed neural networks in modelling industrial systems through case studies in automotive manufacturing: FG Ciampi et al. *International Journal on Interactive Design and Manufacturing (IJDeM)*, 1-11. <https://doi.org/10.1007/s12008-025-02364-w>
- [8] Stanton, I., Munir, K., Ikram, A., & El-Bakry, M. (2023). Predictive maintenance analytics and implementation for aircraft: Challenges and opportunities. *Systems Engineering*, 26(2), 216-237. <https://doi.org/10.1002/sys.21651>
- [9] Xiong, M., Wang, H., Fu, Q., & Xu, Y. (2021). Digital twin-driven aero-engine intelligent predictive maintenance. *The International Journal of Advanced Manufacturing Technology*, 114(11), 3751-3761. <https://doi.org/10.1007/s00170-021-06976-w>
- [10] Huber, L. G., Palmé, T., & Chao, M. A. (2023, June). Physics-informed machine learning for predictive maintenance: Applied use-cases. In *2023 10th IEEE Swiss Conference on Data Science (SDS)* (pp. 66-72). IEEE. <https://doi.org/10.1109/SDS57534.2023.00016>
- [11] Hu, Y., Guo, R., Wang, H., Zhao, R., Ning, R., Huang, Z., ... & Zhang, H. (2025). Gear-fault monitoring and digital twin demonstration of aircraft engine based on piezoelectric vibration sensor for engine health management. *Nano Energy*, 133, 110448. <https://doi.org/10.1016/j.nanoen.2024.110448>

- [12] Huang, B., & Wang, J. (2022). Applications of physics-informed neural networks in power systems-a review. *IEEE Transactions on Power Systems*, 38(1), 572-588. <https://doi.org/10.1109/TPWRS.2022.3162473>
- [13] Zhang, R., Wang, F., Cai, J., Wang, Y., Guo, H., & Zheng, J. (2022). Digital twin and its applications: A survey. *The International Journal of Advanced Manufacturing Technology*, 123(11), 4123-4136. <https://doi.org/10.1007/s00170-022-10445-3>
- [14] Huang, D., Bai, R., Zhao, S., Wen, P., Wang, S., & Chen, S. (2020, June). Bayesian neural network based method of remaining useful life prediction and uncertainty quantification for aircraft engine. In *2020 IEEE International Conference on Prognostics and Health Management (ICPHM)* (pp. 1-8). IEEE. <https://doi.org/10.1109/ICPHM49022.2020.9187044>
- [15] Wang, C., Fan, I. S., & Plastropoulos, A. (2025, March). Digital Twin-Based Health Management for Complex Aircraft Systems: Case Studies and Applications. In *2025 IEEE Wireless Communications and Networking Conference (WCNC)* (pp. 1-6). IEEE. <https://doi.org/10.1109/WCNC61545.2025.10978545>
- [16] Rodrigues, L. R., Gomes, J. P., Ferri, F. A., Medeiros, I. P., Galvao, R. K., & Júnior, C. L. N. (2014). Use of PHM information and system architecture for optimized aircraft maintenance planning. *IEEE Systems Journal*, 9(4), 1197-1207. <https://doi.org/10.1109/JSYST.2014.2343752>
- [17] Afshari, S. S., Jonnadula, B., Xu, X., Liang, X., & Yang, Z. (2022). Jet engine optimal preventive maintenance scheduling using Golden Section search and genetic algorithm. *Journal of Prognostics and Health Management*, 2(1), 45-71. <https://doi.org/10.22215/jphm.v2i1.3321>
- [18] Le, M. Q., Capsal, J. F., Lallart, M., Hebrard, Y., Van Der Ham, A., Reffe, N., ... & Cottinet, P. J. (2015). Review on energy harvesting for structural health monitoring in aeronautical applications. *Progress in Aerospace Sciences*, 79, 147-157. <https://doi.org/10.1016/j.paerosci.2015.10.001>
- [19] Yin, K., Shen, Y., Chen, Y., & Zhang, H. (2025). Multi-Source Information Fusion Diagnosis Method for Aero Engine. *Applied Sciences*, 15(9), 5083. <https://doi.org/10.3390/app15095083>
- [20] Hassani, S., & Dackermann, U. (2023). A systematic review of advanced sensor technologies for non-destructive testing and structural health monitoring. *Sensors*, 23(4), 2204. <https://doi.org/10.3390/s23042204>
- [21] Xiong, J., Fink, O., Zhou, J., & Ma, Y. (2023). Controlled physics-informed data generation for deep learning-based remaining useful life prediction under unseen operation conditions. *Mechanical Systems and Signal Processing*, 197, 110359. <https://doi.org/10.1016/j.ymsp.2023.110359>
- [22] Sadeghi, A., Bellavista, P., Song, W., & Yazdani-Asrami, M. (2024). Digital twins for condition and fleet monitoring of aircraft: toward more-intelligent electrified aviation systems. *IEEE Access*, 12, 99806-99832. <https://doi.org/10.1109/ACCESS.2024.3371902>