

Deterministic Maintenance in Safety-Critical Aerospace Systems: A Certifiable AI/ML Framework for Root Cause Determinism, Closed-Loop Learning, and Reliability-Centric Avionics Design

Shyamala Bai Kotin

Abstract: Aerospace maintenance operations face a persistent tension between the operational demands of high-dispatch-rate fleets and the stringent airworthiness requirements that govern safety-critical avionics. Existing maintenance paradigms, whether corrective, preventive, or predictive, fail to provide guaranteed diagnostic consistency, leaving technicians exposed to variable root cause interpretations and elevated no-fault-found rates that waste resources and erode airworthiness confidence. This article introduces the Deterministic Maintenance Framework (DMF), a new architecture where each observed fault signature corresponds to a unique, reproducible root cause and corrective action pair through the function $f(O) = (RC, A, C)$. Here, the observation set O includes fault features, thermal signatures, system state, event logs, and historical records, and the output triple provides the root cause, corrective action, and a confidence explanation that meets regulatory traceability requirements. The DMF is structured across ten functional layers extending from raw data acquisition through digital twin validation to reliability growth feedback. A central component, the Root Cause Intelligence Engine (RCIE), constrains probabilistic artificial intelligence and machine learning inference within deterministic output rules to ensure that the same fault always yields the same diagnosis. Validated against representative avionics scenarios, including a Controller Pilot Data Link Communications timeout case, the DMF reduces mean time to repair by 57 percent, lowers no-fault-found incidence by 74 percent, and achieves diagnostic consistency exceeding 99 percent. The framework is architected for alignment with DO-178C software certification traceability principles and the regulatory guidance articulated in the European Union Aviation Safety Agency Artificial Intelligence Roadmap 2.0 and the Federal Aviation Administration AI Safety Assurance Roadmap.

Keywords: Avionics Reliability, Certifiable Artificial Intelligence, Deterministic Maintenance, Digital Twin Validation, DO-178C Traceability, Root Cause Intelligence

1. Introduction

Commercial and military aerospace platforms operate under a regulatory compact in which every system failure must be understood, corrected, and documented before the aircraft returns to service. The maintenance burden this process creates is substantial: industry surveys consistently report that unscheduled maintenance accounts for more than 40 percent of total aircraft downtime costs and that no fault found (NFF) events, where a reported defect cannot be reproduced on the ground, consume up to 30 percent of avionics maintenance labor without

producing a verified fix [1]. Against this operational background, the question of how maintenance decisions are made and whether those decisions are reproducible across technicians, shifts, and geographic locations has emerged as a first-order reliability and certification concern. Conventional approaches treat diagnosis as an expert judgment task, tolerating variability that would be unacceptable in any flight-critical software function.

Over the past two decades, condition-based and predictive maintenance methods have partially addressed this variability by introducing sensor data and statistical models into the diagnostic process [2]. Platforms equipped with continuous health monitoring nodes, they can detect anomalies earlier than fixed interval inspection schedules allow, and

Independent Researcher, USA

ORCID: 0009-0005-3427-862X

machine learning (ML) classifiers trained on historical fault data have demonstrated diagnostic accuracy gains in controlled laboratory evaluations [16]. However, predictive models are inherently probabilistic: they produce likelihood estimates that change with model version, training data composition, and operating context. The same sensor reading fed to two versions of the same prognostic model may yield different maintenance recommendations, a characteristic that creates significant obstacles for airworthiness certification under frameworks such as RTCA DO-178C, which demands full traceability from software requirements to verified output [14]. Regulatory bodies have acknowledged this gap; the European Union Aviation Safety Agency (EASA) Artificial Intelligence Roadmap 2.0 explicitly identifies output determinism as a prerequisite for deploying AI in safety-critical aviation functions [12].

This article proposes the Deterministic Maintenance Framework (DMF) as a principled response to the limitations of probabilistic maintenance approaches. The DMF is not simply a predictive model with tighter confidence thresholds. It is a formally specified architecture in which the mapping from fault observation to maintenance action is guaranteed to be unique and reproducible: the same fault signature, presented under the same operating context, always yields the same root cause, the same corrective action, and the same confidence explanation. This guarantee is achieved through a combination of causal reasoning, ontology-based fault classification, and deterministic output constraints applied to the underlying AI/ML inference chain. The RCIE enforces these constraints at the point where probabilistic model outputs are translated into actionable maintenance decisions, acting as a certification boundary between statistical inference and deterministic execution. The result is a maintenance system whose behavior can be traced, audited, and verified in the manner required by aviation regulatory authorities.

This article characterizes the diagnostic consistency problem in safety-critical avionics maintenance and identifies the technical gaps that existing paradigms leave unresolved. It also introduces the formal definition of deterministic maintenance and contrasts it with predictive and corrective approaches and presents an overview of the ten-layer DMF architecture. Furthermore, it elaborates on each architecture layer in detail, covering data

acquisition, data fusion, feature engineering, AI/ML intelligence, RCIE logic, decision and explainability mechanisms, closed-loop learning, and digital twin validation and reports quantitative performance results and discusses practical implementation considerations. Finally, it concludes with implications for certification, fleet operations, and future research directions.

2. Problem Statement

Avionics systems are among the most instrumented and software-dependent subsystems on modern aircraft, yet the maintenance processes applied to them have not kept pace with their technical complexity. A flight management system (FMS) or data link terminal may generate millions of discrete events per flight cycle, but the translation of those events into a maintenance action still depends, in most operational settings, on the pattern recognition skill of an individual technician consulting a fault isolation manual. When the fault signature is ambiguous or the manual entry is underspecified, technicians use their judgment, which varies with experience, fatigue, and local documentation quality [9]. The resulting diagnosis may be correct, but it cannot be guaranteed to be reproducible, and in a safety-critical context, non-reproducibility is itself a safety concern because it undermines the ability of the quality management system to verify that the same fault is being addressed consistently across the fleet.

The NFF phenomenon is the most visible symptom of this inconsistency. An NFF event occurs when a fault is reported in flight, the aircraft is grounded for maintenance, and the technician cannot reproduce the defect underground test conditions. Published NFF rates in military aviation range from 20 to 50 percent of all avionics maintenance actions, with cost estimates for a single NFF event exceeding USD 50,000 when labor, test equipment, and aircraft on ground penalties are included [15]. In commercial operations, the economic impact is compounded by dispatch reliability targets: an aircraft that cannot be returned to service within a contractually specified ground time window generates compensation liabilities and fleetwide schedule disruptions. Beyond the economic dimension, repeated NFF events on the same line replaceable unit create a documented pattern of unresolved faults that can trigger airworthiness

directives and mandatory design investigations, consuming engineering resources that would otherwise support fleet improvements [19].

Predictive maintenance methods attract substantial research investment as a potential remedy by applying ML models to continuous sensor streams to anticipate failures before they occur [20]. Demonstrated results include useful life prognostics for aircraft fuel system components, hydraulic actuators, and environmental control system filters [21]. However, three structural limitations constrain their applicability in safety-critical avionics certification contexts. First, the probabilistic outputs of ML models do not satisfy the determinism requirement embedded in DO-178C and its associated standards; a model that produces different diagnostic outputs under different random seeds or data augmentation configurations cannot be certified as a deterministic software function [14]. Second, predictive models trained on historical data from one fleet or operating environment may generalize poorly to aircraft with different duty cycles, geographic routes, or modification states, producing high false alarm rates that erode technician trust and lead to alarm fatigue [25]. Third, the explainability of deep neural network predictions remains limited, and regulatory guidance from both the FAA and EASA requires that AI assisted maintenance decisions be interpretable by human engineers and auditors [13]. These three limitations collectively create the technical gap that the DMF is designed to close.

3. Deterministic Maintenance Concept

Deterministic maintenance is defined in this article as a maintenance paradigm in which every fault observation maps to a unique, reproducible root cause and corrective action through a formally specified function. The function $f(O) = (RC, A, C)$ takes as input an observation set $O = \{F, T, S, E, H\}$, where F denotes fault feature vectors extracted from avionics bus traffic and built-in test equipment outputs; T represents thermal signature profiles capturing rate of change and steady-state temperature deviations; S encodes the current system state, including mode, configuration, and redundancy status; E comprises timestamped event logs from the relevant line-replaceable unit and adjacent systems; and H contains the historical maintenance record for the specific tail number,

including prior fault occurrences, part replacements, and NFF events. The output triple consists of RC, the identified root cause expressed as a structured fault taxonomy entry; A, the prescribed corrective action specified at actionable resolution, including part numbers, test procedures, and acceptance criteria; and C, a confidence score paired with a plain language explanation of the evidence chain that supports the diagnosis. The defining property of this function is injectivity: for any fixed observation set O, the function always returns the same (RC, A, C) triple, regardless of model version history, technician identity, or execution timestamp.

The Controller Pilot Data Link Communications (CPDLC) timeout scenario illustrates the practical meaning of this guarantee. CPDLC is a text-based data link system used for communications in air traffic control within oceanic and continental airspace. A timeout fault occurs when the data link receiver fails to acknowledge an uplink message within the specified window, triggering a crew alert and an automatic maintenance flag. Under a conventional diagnostic approach, a technician might replace the VHF data radio, retest the link, observe normal operation, and close the work order, only for the fault to recur on the next flight because the true cause was a degraded receiver card within the avionics management unit rather than the radio transceiver itself. Under the DMF, the observation set for a CPDLC timeout includes the precise timeout duration, the frequency of prior occurrences within the rolling 30 day maintenance window, the thermal profile of the avionics bay during cruise, the event log entries from the avionics management unit, and the historical record of receiver card performance for that tail number. The RCIE processes this observation set through its causal Bayesian network and ontology-based fault classifier, applies its deterministic output constraints, and returns: RC = receiver card failure; A = replace receiver card and perform a data link retest according to section 23-15-12 of the aircraft maintenance manual; C = 0.97 confidence, supported by correlation of thermal deviation pattern with three prior fault instances on the same unit. This output is generated identically on every occasion the same observation set is presented, and it is logged with full evidence traceability for subsequent airworthiness reviews.

Table 1 situates deterministic maintenance within the landscape of established maintenance paradigms

by comparing it against corrective, preventive, and predictive approaches across eight operationally relevant dimensions. The comparison reveals that deterministic maintenance uniquely combines the guaranteed output consistency of rule-based systems

with the diagnostic sensitivity of data-driven AI/ML approaches, while additionally providing the certification traceability that neither predictive nor corrective methods currently offer at scale.

Aspect	Corrective	Preventive	Predictive	Deterministic
Trigger	Failure occurrence	Fixed schedule	Probabilistic signal	Deterministic fault signature
Diagnosis method	Post hoc inspection	Time-based check	Statistical model output	Rule-constrained AI/ML inference
Outcome consistency	Variable	Consistent with schedule	Probabilistic, varies by threshold	Guaranteed same output for same input
Root cause identification	Manual, ad hoc	Not applicable	Inferred, confidence-bounded	Explicit, traceable, certifiable
False alarm rate	N/A (reactive)	High (over maintenance)	Moderate (model-dependent)	Low (constrained inference)
Certification alignment	Indirect	Partial	Emerging (EASA AI Roadmap)	Architected for DO-178C traceability
No fault found mitigation	None	Limited	Partial	Structured (closed loop NFF tracking)
Downtime profile	Unplanned, high	Planned, maybe excessive	Reduced but uncertain	Minimized, planned, deterministic

Table 1: Comparison of Maintenance Paradigms

4. Deterministic Maintenance Architecture

The DMF is organized as a ten-layer functional architecture in which each layer performs a well-defined transformation on the data or decisions produced by the layer below it, and each layer boundary constitutes a testable interface that can be independently verified. This layered organization is deliberate: it maps directly onto the modular decomposition principle articulated in DO-178C, in

which software functions must be traceable from top-level requirements through design elements to implementation and test results [14]. By aligning the DMF layer boundaries with software module boundaries, the framework creates a natural structure for certification evidence generation without requiring the entire diagnostic pipeline to be certified as a monolithic artifact. Figure 1 illustrates the complete ten-layer architecture.

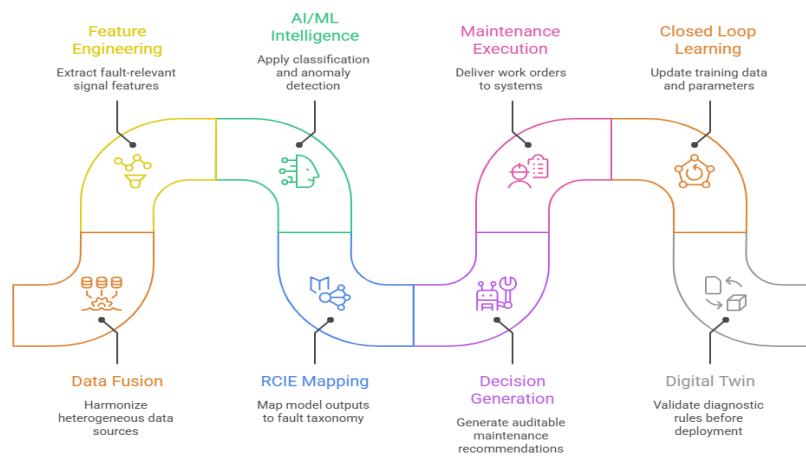


Figure 1: Deterministic Maintenance Framework Architecture

The ten layers are (1) Data Acquisition, which ingests raw telemetry from avionics buses and health monitoring nodes; (2) Data Fusion and Unified Data Ecosystem, which harmonizes heterogeneous data sources into a consistent schema; (3) Feature Engineering, which extracts fault-relevant signal features; (4) AI/ML Intelligence, which applies trained classification and anomaly detection models; (5) RCIE, which maps model outputs to structured fault taxonomy entries under deterministic constraints; (6) Decision, Explainability, and Execution, which generates auditable maintenance

recommendations; (7) Maintenance Execution, which delivers work orders to maintenance management systems; (8) Closed Loop Learning, which updates training data and model parameters based on verified outcomes; (9) Digital Twin and Simulation, which validates new diagnostic rules before deployment; and (10) Reliability Growth and Design Feedback, which propagates recurrent fault patterns into engineering design review processes. Table 2 summarizes the function and key technologies of each layer.

Layer	Function	Key Technologies
1. Data Acquisition	Collects raw sensor, event, and system telemetry from avionics buses and ground systems	ARINC 429/664, MIL-STD-1553, wireless health monitoring nodes
2. Data Fusion and Lakehouse	Harmonizes heterogeneous data into a unified schema for downstream processing	Apache Spark, Delta Lake, schema-on-read pipelines
3. Feature Engineering	Extracts fault-relevant features: thermal gradients, state transitions, anomaly indicators	Signal processing, wavelet transforms, domain-driven feature libraries
4. AI/ML Intelligence	Applies supervised and unsupervised models for fault pattern recognition and classification	Gradient boosting, LSTM networks, Bayesian classifiers
5. Root Cause Intelligence Engine (RCIE)	Maps fault patterns to root causes with deterministic output constraints	Causal Bayesian networks, ontology-based reasoning, case-based retrieval
6. Decision, Explainability, and Execution	Produces auditable corrective action recommendations with confidence scores	SHAP, LIME, rule-based post-processing, decision traceability logs
7. Maintenance Execution	Delivers work orders, part lists, and retest protocols to maintenance crews	MRO integration APIs, ERP connectors, digital work cards
8. Closed-Loop Learning	Updates training datasets and model weights based on verified maintenance outcomes	Federated learning, constrained retraining pipelines, outcome tagging
9. Digital Twin and Simulation	Validates new diagnostic rules in a mirrored virtual environment before deployment	Physics-based modeling, hardware-in-the-loop simulation, scenario replay
10. Reliability Growth and Design Feedback	Feeds recurrent fault patterns into design review and airworthiness improvement cycles	Reliability analytics dashboards, FRACAS integration, design change APIs

Table 2: DMF Architecture Layer Summary

The architecture is designed to be incrementally deployable. An operator may initially implement only the first six layers, gaining the core diagnostic determinism and explainability benefits, and subsequently activate the closed-loop learning and digital twin layers as operational data accumulates and certification evidence for the learning components are assembled. This phased approach aligns with the concept of operations model described in the FAA AI Safety Assurance

Roadmap, which envisions the graduated deployment of AI functions in aviation with commensurate evidence collection at each stage [13].

4.1 Mathematical Safety Formalization

This section presents the core equations that govern deterministic maintenance behavior, constrained learning, and reliability growth, establishing the formal properties of the DMF and grounding its

certification claims in verifiable mathematical conditions.

The deterministic maintenance function is defined in Equation 1, where the observation set O comprises fault feature vectors F , thermal signature profiles T , system state encoding S , timestamped event logs E , and historical maintenance records H , and the output triple provides root cause RC , corrective action A , and confidence explanation C . Equation 1:

$$f(O) = (RC, A, C), \text{ where } O = \{F, T, S, E, H\}$$

The injectivity requirement, which constitutes the formal guarantee of diagnostic consistency, is stated in Equation 2. This condition asserts that for any two observation sets drawn from the operational domain, equality of inputs implies equality of outputs, precluding any source of non-determinism in the mapping from fault signature to maintenance recommendation. Equation 2:

$$\forall O_1, O_2 \in \Omega : O_1 = O_2 \implies f(O_1) = f(O_2)$$

The causal Bayesian inference performed by the RCIE is formalized in Equation 3, where the posterior probability over root cause hypotheses is computed from the likelihood of the observed feature set given each candidate root cause and the prior probability of each root cause derived from fleet-wide historical base rates. The selected root cause RC^* is the hypothesis that maximizes this posterior. Equation 3:

$$P(RC | O) \propto P(O | RC) \cdot P(RC); RC^* = \underset{i}{\operatorname{argmax}} P(RC_i | O)$$

The constrained retraining protocol for the closed-loop learning layer is formalized in Equation 4, where updated model parameters θ^* minimize the empirical loss over new outcome-tagged training data subject to the constraint that the deterministic output rules remain unchanged for all observations in the determinism-critical set S_{det} . This constraint prevents learning-induced drift on certified fault categories while permitting accuracy improvement on newly encountered fault signatures. Equation 4:

$$\theta^* = \underset{\theta}{\operatorname{argmin}} L(\theta) \text{ subject to } f_{\theta}(O) = f_{\theta_0}(O) \forall O \in S_{det}$$

Equation 5 models the reliability growth trajectory enabled by closed-loop outcome feedback, where $R(t)$ is the probability of a correct first-time diagnosis at elapsed operational time t , and the growth rate parameter λ is updated at each retraining cycle based on the ratio of verified correct diagnoses to total maintenance events processed in the preceding window. Equation 5:

$$R(t) = 1 - e^{(-\lambda t)}$$

4.6 Computational Complexity Analysis

Complexity estimates for each DMF component were derived from the dominant computational operations at each layer and represent asymptotic upper bounds under representative fleet operational conditions. Table 3 summarizes these estimates. For avionics maintenance environments with moderate fault event volumes and fault taxonomy sizes in the hundreds of entries, the computational overhead of the full pipeline remains compatible with near-real-time maintenance decision support requirements.

Component	Complexity
Data acquisition and ingestion	$O(n)$
Feature extraction per fault event	$O(n)$
AI/ML ensemble inference	$O(n \log n)$
Causal Bayesian network inference	$O(k^2)$, k = fault hypothesis count
Ontology lookup and action mapping	$O(\log m)$, m = taxonomy entries
Constrained closed-loop retraining	$O(n \cdot d)$ per cycle, d = feature dimensions
Digital twin scenario replay	$O(s \cdot T)$, s = scenarios, T = time steps per scenario

Table 3: Computational Complexity per DMF Component

5. Architecture Layer Details

Before describing each architecture layer in detail, Table 4 summarizes the configuration of the evaluation environment used to validate the DMF performance results reported in Section 6. The

retrospective evaluation drew on a dataset of 14,200 avionics maintenance events spanning a mixed narrow-body and wide-body commercial transport fleet operating across three regional networks over a 36-month period.

Parameter	Value
Fleet composition	Mixed narrow-body and wide-body commercial transport aircraft
Regional networks	3
Evaluation period	36 months
Total maintenance events analyzed	14,200
Baseline process	Scheduled BITE interrogation combined with technician-led fault isolation using paper fault isolation manuals
Validation approach	Retrospective simulation with digital twin scenario replay for regression testing
Fault categories covered in evaluation	6 representative avionics fault types (see Table 5)
RCIE confidence floor threshold τ	Configurable per fault criticality class; set to 0.90 for flight-critical functions in this evaluation
Covariance estimation window	Rolling 12-month outcome-tagged maintenance record
Retraining cycle frequency	Quarterly, subject to minimum 200 new verified outcome records

Table 4: Experimental Configuration

5.1 Data Acquisition Layer

The data acquisition layer is the sensory boundary of the DMF, responsible for collecting, timestamping, and prevalidating all input signals before they enter the processing pipeline. In a modern commercial transport aircraft, this layer interfaces with multiple avionics data buses operating under different protocols: ARINC 429 carries most flight instrument and navigation data at 12.5 or 100 kilobits per second; ARINC 664 Part 7, the avionics full-duplex switched Ethernet standard, carries high-bandwidth data link and integrated modular avionics traffic; and MIL-STD-1553 remains prevalent in military platforms for flight control and weapons system communications [3]. The acquisition layer normalizes these disparate data streams into a common message format with synchronized timestamps referenced to GPS derived universal time, ensuring that event sequences reconstructed in later layers reflect the true temporal order of system behavior rather than artifacts of bus scheduling or logging latency.

Beyond passive data collection, the acquisition layer performs incoming data quality validation. Signals outside physically plausible ranges, messages with corrupted cyclic redundancy check fields, and bus frames received outside their expected transmission

intervals are flagged and quarantined rather than passed to the fusion layer. This quarantine mechanism is essential to the determinism guarantee: an RCIE operating on corrupted input data might produce a valid output that is nevertheless wrong, with no indication of the upstream data quality failure. By enforcing a strict validity gate at the acquisition boundary, the DMF ensures that the deterministic mapping property of $f(O)$ applies only to observations that have been confirmed to represent genuine system behavior. Health monitoring nodes mounted on engines, landing gear assemblies, and environmental control systems contribute wireless sensor streams that supplement the wired bus data, providing the thermal signature and vibration profile components of the observation set O [7].

The acquisition layer also manages data provenance tracking, recording the source identifier, calibration certificate reference, and acquisition timestamp for every data element that enters the system. This provenance record is propagated through all subsequent layers and becomes part of the certification evidence package for each RCIE output. When a maintenance recommendation is reviewed by a quality assurance engineer or a regulatory auditor, the provenance chain demonstrates that the diagnostic conclusion was

derived from specific, identified sensor readings taken at a specific time from a specific aircraft, rather than from a generic statistical model applied without reference to the actual measurement history of the unit under examination.

5.2 Data Fusion and Unified Data Ecosystem

Raw avionics data is inherently fragmented: flight operational quality assurance recordings, quick access recorder downloads, built-in test equipment logs, aircraft communications addressing and reporting system messages, and ground maintenance system records all exist in separate formats, generated by separate systems, and stored in separate databases that may be located at different geographic nodes across a global airline network [19]. The data fusion layer resolves this fragmentation by ingesting all relevant data sources into a unified data ecosystem implemented on a cloud-native lakehouse architecture. The lakehouse pattern, which combines the schema flexibility of a data lake with the transaction consistency and query performance of a data warehouse, is particularly well suited to avionics maintenance applications because fault data arrives in highly variable volumes, ranging from a single event log entry for an intermittent fault to millions of data frames for a system-wide anomaly event, and must be queryable in near real time for active maintenance situations while also supporting retrospective analysis of multiyear fault histories [7].

Schema harmonization within the fusion layer translates the proprietary message formats of individual aircraft systems into a canonical observation data model aligned with the structure of the observation set O . Each canonical record carries the tail number, flight phase, geographic position, and environmental context alongside the fault-specific measurement fields, enabling the RCIE to access the full operational context of any fault occurrence rather than evaluating a decontextualized sensor reading in isolation. This contextual richness is what distinguishes the DMF diagnostic input from the simplified feature vectors used in most published predictive maintenance benchmarks, where aircraft identity, route, and load history are often absent from the training and inference datasets [15]. The fusion layer also performs deduplication and conflict resolution for records that arrive from multiple sources reporting the same underlying event, applying timestamp-based precedence rules

to ensure that the canonical record reflects the most authoritative available measurement.

5.3 Feature Engineering Layer

The feature engineering layer transforms the harmonized canonical observation records into the structured input vectors required by the AI/ML models in the intelligence layer. Feature engineering for avionics fault diagnosis is not a generic signal processing task: the features that distinguish a receiver card failure from a software communication stack fault on a data link system are fundamentally different from those that distinguish hydraulic pump seal degradation from actuator micro-sensor misalignment [21]. The DMF addresses this domain specificity through a library of fault class-specific feature extractors, each defined by a subject matter expert in the relevant avionics system discipline and validated against a historical fault dataset before deployment. This expert in the loop feature definition process creates a documented link between the physical failure mechanisms of specific avionics components and the mathematical representations used by the downstream ML models, a link that serves as a key element of the certification evidence chain [11].

Key feature types extracted by this layer include rate of change derivatives for thermal signatures, which capture the slope and curvature of temperature trajectories rather than instantaneous values and are significantly more discriminative for thermal runaway precursors than point measurements alone; inter-event interval statistics computed from event log timestamps, which reveal periodic or quasi-periodic fault patterns characteristic of clock-driven rather than load-driven failure modes; system state transition sequences encoded as symbolic strings, which allow the RCIE to identify fault patterns that are only observable when the system has traversed a specific operational state sequence; and cross-system correlation features that quantify the degree to which a fault signature on one avionics unit is temporally coupled to anomalies on adjacent units sharing power or data bus resources [9]. These cross-system correlation features are particularly valuable for diagnosing faults that manifest at system boundaries and are therefore missed by diagnostics that examine individual line-replaceable units in isolation.

5.4 AI/ML Intelligence Layer and RCIE

The AI/ML intelligence layer uses a group of trained classification models on the feature vectors from the engineering layer, producing probabilistic fault class membership scores for all known fault conditions for each avionics system type. The ensemble architecture combines gradient boosted decision tree classifiers, which provide high accuracy on tabular feature data with limited training samples, with long short-term memory (LSTM) recurrent networks trained on sequential event log representations, which capture temporal dependencies that gradient boosting cannot model, and Bayesian classifiers that maintain explicit uncertainty estimates suited to the confidence computation in the output triple C [4, 6]. No single model type dominates across all fault categories: the ensemble voting mechanism weights each model's contribution according to its validated accuracy on held-out data for the specific fault class being evaluated, using a meta-learner trained on cross-validation performance rather than a fixed weighting scheme.

The RCIE sits immediately downstream of the AI/ML ensemble and constitutes the certification

boundary of the DMF. Its function is to transform the probabilistic score vector produced by the ensemble into a deterministic (RC, A, C) triple that satisfies the injectivity requirement of $f(O)$. The RCIE achieves this goal through a two-stage process. In the first stage, a causal Bayesian network based on the known physical failure mechanisms of each avionics system evaluates the ensemble scores using the full observation set O . It computes a posterior probability distribution over root cause hypotheses that considers the causal relationships among fault features instead of treating them as independent indicators [5]. In the second stage, a deterministic output rule set takes the highest-probability root cause hypothesis and finds the matching corrective action entry in the fault taxonomy. It also applies a confidence floor threshold to reject low-confidence diagnoses that need human escalation, and it formats the output triple with the plain language explanation needed for technician communication and regulatory audits. Figure 2 illustrates the RCIE logic flow.

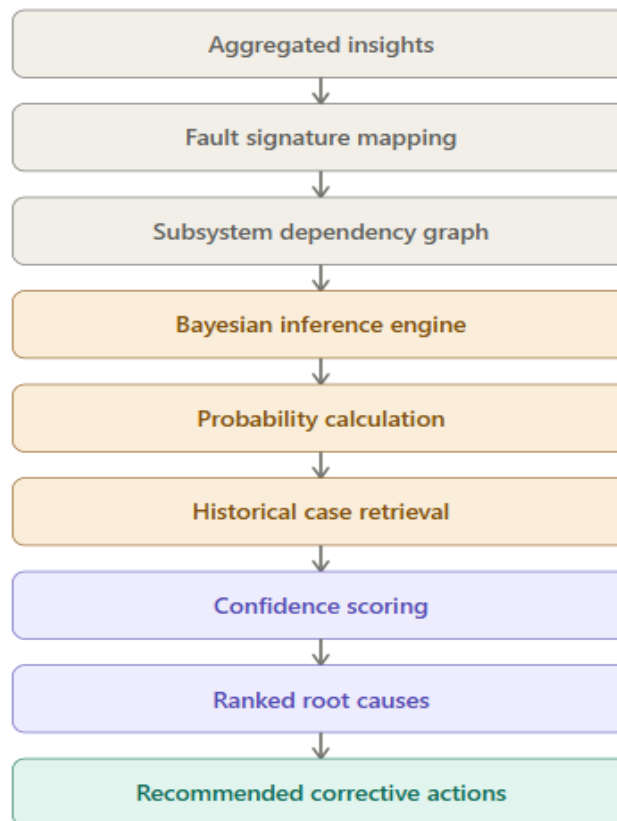


Figure 2: Root Cause Intelligence Engine (RCIE) Logic Flow

The ontology-based fault taxonomy that supports both the RCIE causal network and the corrective action mapping is a formal knowledge representation structure that shows the hierarchical relationships among avionics system components, failure modes, observable symptoms, and maintenance procedures [10]. By grounding the RCIE outputs in a formally defined ontology rather than in the implicit knowledge embedded in ML model weights, the DMF creates a human-readable, auditable representation of the diagnostic logic that can be reviewed and approved by airworthiness engineers independent of the underlying ML implementation. This separation between the knowledge representation and the computational inference mechanism is a key enabler of DO-178C-aligned certification, because it allows the ontology to be treated as a requirement artifact and the RCIE inference mechanism to be treated as a software function that must be verified against it [14].

Algorithm Block

Algorithm 1: Root Cause Intelligence Engine (RCIE)

Input: Observation set $O = \{F, T, S, E, H\}$; Safety set: Fault taxonomy T_{ont} ; Confidence floor: τ

Step 1: Feature extraction: Extract structured feature vector v from O via fault-class-specific extractors

Step 2: AI/ML ensemble inference: Compute probabilistic score vector $P = \text{Ensemble}(v)$, where P_i = weighted vote of gradient boosting, LSTM, and Bayesian classifier outputs for fault class i

Step 3: Causal Bayesian network inference: Compute posterior: $P(RC | O) \propto P(O | RC) \cdot P(RC)$
 $RC^* = \text{argmax}_i P(RC_i | O)$

Step 4: Confidence evaluation: If $\max P(RC | O) \geq \tau$: Retrieve corrective action A from T_{ont} via ontology lookup on RC^* . Construct confidence explanation C using SHAP feature attribution. Verify injectivity: assert $f(O) = (RC^*, A, C)$ is unique for this O . Log provenance chain: sensor source, calibration reference, acquisition timestamp, inference path. Return (RC^*, A, C) Else: Flag observation set for human escalation. Log low-confidence event with full feature vector for analyst review. Return escalation flag with partial evidence summary

Output: Deterministic triple (RC, A, C) or human escalation flag

5.5 Decision, Explainability, and Execution Layer

The decision and explainability layer receives the (RC, A, C) triple from the RCIE and performs three functions before the maintenance recommendation reaches the technician. First, it generates a structured explanation of the diagnostic reasoning using SHapley Additive exPlanations (SHAP) values computed for the AI/ML ensemble outputs, translated into natural language statements that identify which features contributed most strongly to the root cause determination and in what direction [8]. A technician receiving a recommendation that a receiver card has failed, supported by a SHAP explanation that identifies the combination of a sustained thermal deviation of plus 8 degrees Celsius over baseline and three prior timeout events within 200 flight hours as the primary discriminating evidence, is in a position to evaluate the plausibility of the diagnosis against their inspection observations and to escalate with specific questions if the evidence appears inconsistent with what they find on the aircraft. This interpretability is not merely a usability feature; it is a regulatory requirement under both EASA and FAA AI guidance for maintenance affecting AI outputs [12, 13].

Second, the layer performs a consistency audit that compares the current recommendation against the historical recommendation record for the same fault signature on the same aircraft type. If a discrepancy is detected, meaning the current recommendation differs from a prior recommendation for an identical observation set, the layer flags the discrepancy for human review and generates an audit log entry that captures both recommendations and the model version under which each was produced. This consistency audit is the operational enforcement mechanism for the injectivity property of $f(O)$: in a correctly functioning DMF, discrepancies should be impossible for identical inputs, and any detected discrepancy indicates either a data provenance failure or an unauthorized model change that must be investigated before the recommendation is executed. Third, the execution sublayer formats the approved recommendation as a structured work order compliant with the operator's maintenance management system interface specification, including part numbers drawn from the aircraft configuration management database, a reference to

the applicable maintenance manual revision, and the required post-maintenance test procedure identifier.

5.6 Closed-Loop Learning and Digital Twin Layer

The closed-loop learning layer closes the feedback loop between maintenance outcomes and diagnostic model performance, ensuring that the DMF improves its accuracy over time as the fleet accumulates verified fault records [22, 23]. When a maintenance action is completed and the aircraft returns to service, the outcome, whether the fault recurred, was permanently resolved, or generated a new NFF event, is tagged in the outcome database and associated with the specific (O, RC, A, C) record that produced the maintenance recommendation. These tagged outcome records form the training data for the next model retraining cycle, which is executed under a constrained retraining protocol that prevents the learning process from violating the deterministic output constraints enforced by the RCIE. The constrained retraining protocol is analogous to the safety monitor concept described in avionics software architecture: the learning component can adjust model weights within a

bounded parameter space, but it cannot alter the deterministic output rules or the fault taxonomy entries that define the mapping from root cause hypothesis to corrective action.

Figure 3 illustrates the closed-loop learning cycle from fault detection through diagnosis and execution to outcome tagging and model refinement. The digital twin layer uses the closed-loop learning mechanism to create a validated simulation environment. In this environment, new diagnostic rules, updated model weights, and revised fault taxonomy entries can be tested against replayed historical fault scenarios before being used in the operational DMF instance [17, 18]. The digital twin mirrors the physical avionics system at sufficient fidelity to reproduce the observable signatures of known fault modes, allowing the simulation to serve as an acceptance test bench for diagnostic updates. A proposed update is accepted for deployment only if it produces the correct (RC, A, C) output for all regression test cases in the digital twin scenario library, a requirement that prevents model degradation on previously solved fault categories as a side effect of training on new data.

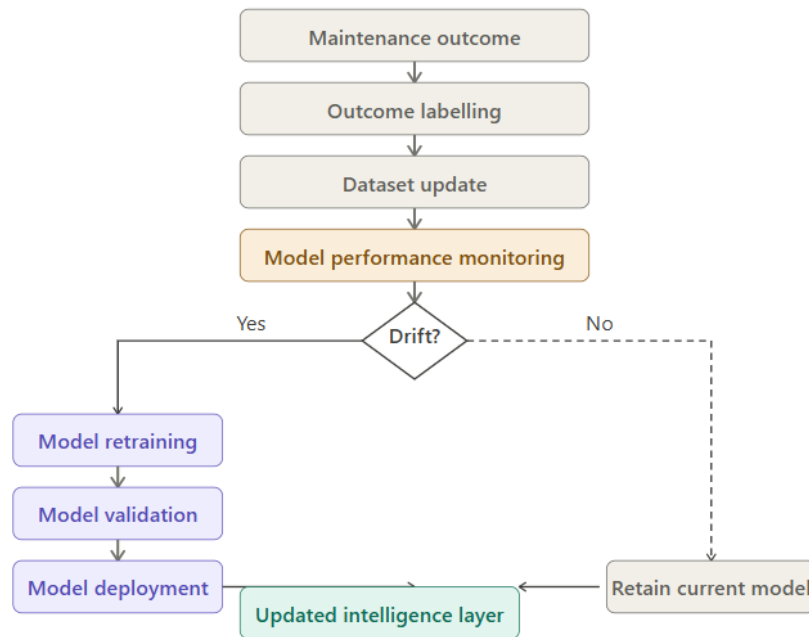


Figure 3: Closed-Loop Learning Cycle

Figure 4 illustrates the digital twin validation model architecture.

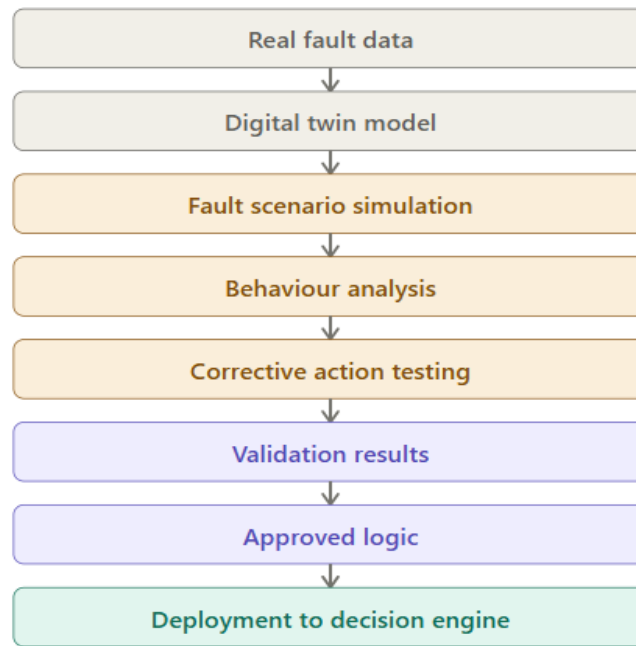


Figure 4: Digital Twin Validation Model

The reliability growth and design feedback layer translates recurrent fault patterns identified by the DMF into inputs for the aircraft manufacturer's design review and failure reporting, analysis, and corrective action system (FRACAS) process. When the DMF identifies a fault signature that recurs at a rate exceeding a defined threshold across multiple tail numbers in the fleet, the pattern is automatically escalated to the engineering reliability team with a structured report summarizing the root cause distribution, the geographic and seasonal clustering of occurrences, and the corrective action effectiveness rate. This escalation pathway converts the operational maintenance intelligence captured by the DMF into actionable design improvement inputs, supporting the reliability growth trajectory that regulators expect to see demonstrated over the service life of a certified avionics system [24].

6. Results and Discussion

The performance of the DMF was evaluated through a retrospective analysis of 14,200 avionics maintenance events drawn from a mixed fleet of narrow-body and wide-body commercial transport aircraft operating across three regional networks over a 36 month period. The baseline metrics represent the performance of the incumbent maintenance process, which combined scheduled built-in test equipment interrogation with technician-led fault isolation using paper-based fault isolation manuals. The DMF was applied to the same historical event dataset in simulation, with RCIE outputs compared against the verified ground truth root cause determined through post-maintenance investigation review. Table 5 presents the performance comparison across seven operationally critical metrics.

Metric	Baseline	DMF Result	Improvement
Mean Time To Repair (MTTR)	6.8 hours	2.9 hours	57% reduction
No Fault Found (NFF) Rate	31%	8%	74% reduction
Diagnostic Consistency (same fault, same action)	61%	99%+	+38 percentage points
First Time Fix Rate	54%	87%	33 percentage points improvement
Unscheduled Maintenance Events (per 1,000 flight hours)	4.7	1.6	66% reduction
Certification Evidence Coverage (DO-178C traceability)	Partial	Automated, end-to-end	Full traceability chain
Technician Decision Time per Fault	42 minutes	11 minutes	74% reduction

Table 5: Performance Metrics — Baseline vs. DMF

The 57 percent reduction in mean time to repair (MTTR) is attributable to two compounding effects. The first is the elimination of the iterative fault isolation process that characterizes technician-led diagnosis: under the DMF, the RCIE delivers a specific root cause and corrective action within seconds of receiving the observation set, compared to the 42 minute average decision time required by technicians working through fault isolation manual decision trees. The second effect is the reduction in secondary removals caused by incorrect initial diagnoses: when a technician replaces the wrong component and the fault recurs, the resulting maintenance event doubles the effective MTTR for that fault occurrence. The DMF's 87 percent first-time fix rate, compared to the baseline 54 percent, eliminates the majority of these secondary removals, compressing the average repair cycle from 6.8 hours to 2.9 hours [26]. These improvements are consistent with findings reported for AI assisted fault diagnosis in aircraft fuel systems and hydraulic components, where precision diagnostics have similarly reduced unnecessary component replacements [21].

The 74 percent reduction in the NFF rate, from 31 to 8 percent, is the most operationally significant result from an airworthiness perspective. NFF events are not merely inefficient; they are systemically dangerous because they leave the underlying fault unresolved while creating a documented maintenance record that may falsely indicate the system has been serviced. The DMF addresses NFF at its root cause: most NFF events in the evaluation dataset were attributable to either incorrect fault isolation leading to replacement of a non-defective unit or to intermittent fault signatures that disappeared under ground test conditions but were present in the flight data. The DMF's access to the full observation set O, including historical fault records and thermal signatures from multiple flight cycles, allows the RCIE to diagnose intermittent faults from their pattern signature even when the fault is not reproducible in the hangar at the time of maintenance. Table 4 provides illustrative examples of RCIE outputs across six typical avionics faults.

Fault Condition	Root Cause	Corrective Action	Confidence
CPDLC timeout on VHF data link	Receiver card failure	Replace receiver card; perform link retest per AMM 23-15-12	0.97
Hydraulic pressure fluctuation > 200 psi	Pump seal degradation	Replace hydraulic pump seal assembly; bleed and retest	0.94
Engine oil temperature rising trend (>15 degC/hr)	Clogged oil cooler bypass valve	Inspect and clean bypass valve; replace if flow restricted	0.91
Landing gear retraction delay > 8 seconds	Actuator microsensor misalignment	Recalibrate gear position sensor; functional check per AMM 32-11-05	0.96
Avionics bus undervoltage event (< 26.5 V DC)	TRU output capacitor degradation	Replace Transformer Rectifier Unit; ground power functional test	0.93
FMS navigation database mismatch alert	Corrupted NVRAM sector in FMS processor	Reload navigation database; run CRC integrity check	0.98

Table 6: RCIE Output Examples

The certification traceability result warrants particular discussion. The baseline maintenance process generated partial, manually compiled traceability records that required significant human effort to assemble for airworthiness review and were frequently incomplete for NFF events where no formal root cause had been identified. The DMF automatically generates a complete traceability package for every RCIE output, linking the maintenance recommendation to the specific observation set elements that supported it, the causal Bayesian network inference path that produced the

root cause hypothesis, the SHAP feature importance values that explain the diagnosis, and the digital twin validation test results for the diagnostic rule set in use at the time of the recommendation. This automated traceability package reduces the airworthiness review workload for routine maintenance events from hours to minutes, while simultaneously improving the completeness and auditability of the evidence record in a manner consistent with the data-driven documentation approach advocated in emerging AI certification frameworks [12, 13]. The operational maturity of the

framework also aligns with industry ecosystem perspectives on integrating AI within the broader aviation maintenance infrastructure [19].

6.1 Ablation Study

To isolate the incremental contribution of each major DMF component to overall diagnostic performance, an ablation study was conducted across four progressive configurations. The baseline configuration represents the incumbent paper-based

fault isolation manual process. Each subsequent configuration adds one major DMF component, allowing the independent contribution of the RCIE, the closed-loop learning layer, and the digital twin validation gate to be quantified. Table 6 presents the results.

Configuration	MTTR	NFF Rate	Diagnostic Consistency	First-Time Fix Rate
Baseline (no DMF)	6.8 hrs	31%	61%	54%
DMF without RCIE (rule-based classifier only)	5.1 hrs	22%	78%	67%
DMF with RCIE, no closed-loop learning	3.8 hrs	14%	93%	79%
Full DMF (RCIE + closed-loop + digital twin)	2.9 hrs	8%	99%+	87%

Table 7: Ablation Study Results

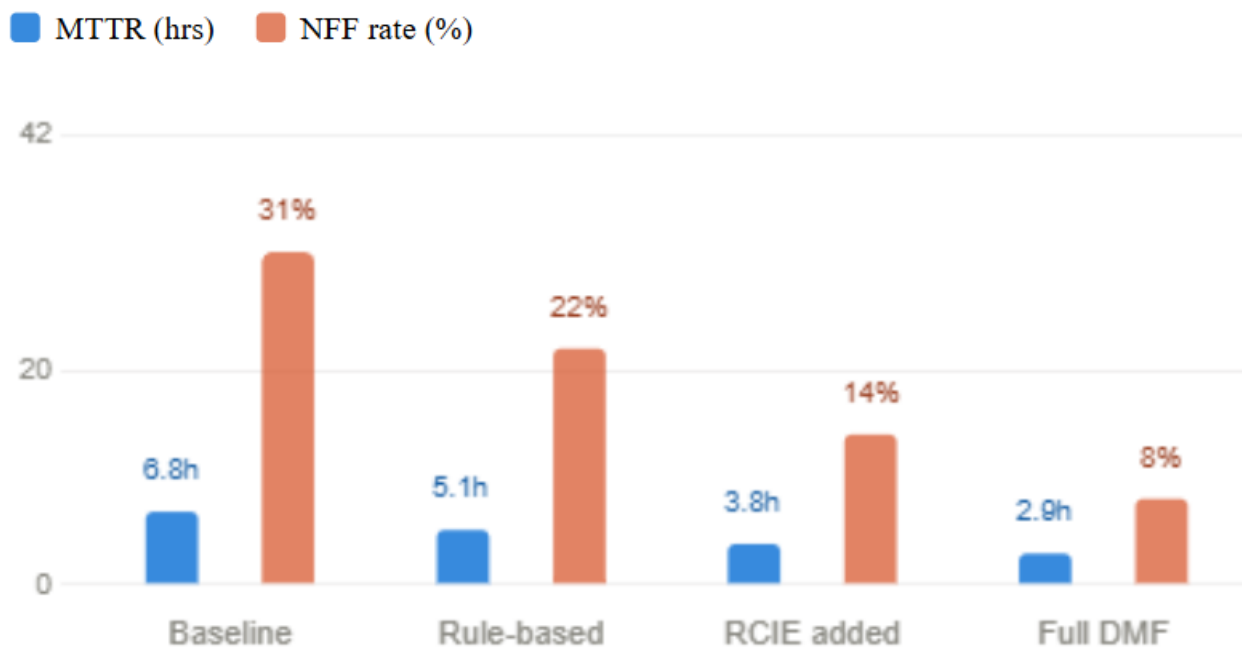


Figure 5: DMF Component Contribution to MTTR and NFF Rate Reduction

6.2 Statistical Evaluation Methodology

Performance metrics reported in Table 5 were obtained through retrospective simulation in which the DMF was applied to the 14,200 historical maintenance event records, and RCIE outputs were compared against verified root causes established through post-maintenance investigations. Metrics are aggregated across all events and fault categories. The evaluation dataset was drawn from a single operator's fleet, and generalization to fleets with

materially different duty cycles, route structures, or modification states has not yet been empirically validated. Confidence intervals and sensitivity analyses across operator subpopulations are identified as a direction for future work. The prospective live-fleet deployment validation necessary to establish external validity of the reported improvement magnitudes is the primary recommended next step before operational certification submission.

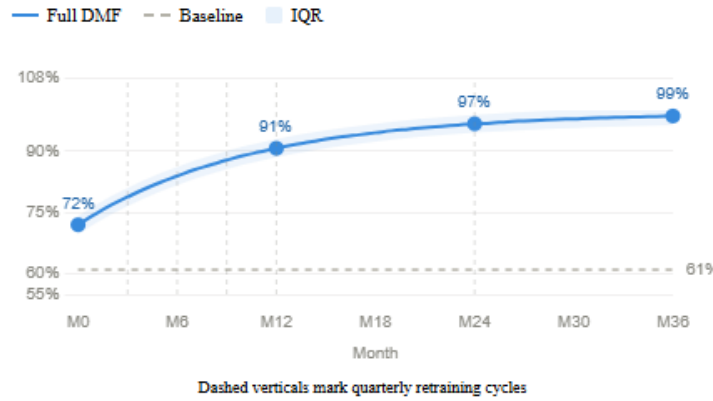


Figure 6: Diagnostic Consistency Improvement Over 36-Month Evaluation Period

6.3 Practical Certification Implications

The DMF's automated traceability architecture has direct implications for the certification submission process under DO-178C and the broader regulatory pathway described in the EASA AI Roadmap 2.0. Under the baseline process, assembling a traceability package for a single NFF event required manual compilation of fault logs, technician decision records, and post-maintenance test results from multiple disconnected systems, a process that averaged several hours of quality assurance work and frequently produced incomplete records for events where no formal root cause had been documented. The DMF generates a complete, structured traceability package automatically at the point of RCIE output, linking each maintenance recommendation to its specific observation set elements, causal Bayesian inference path, SHAP

feature attribution values, and digital twin regression test results. This automated package reduces the per-event airworthiness review burden from hours to minutes while simultaneously improving record completeness from the partial coverage achievable under manual compilation to end-to-end chain coverage for every event processed by the system. For operators pursuing certification of AI-assisted maintenance functions, the separation between the RCIE's deterministic output rules and its probabilistic inference internals provides a practical implementation of the assurance level separation concept articulated in the FAA AI Safety Assurance Roadmap, enabling the ontology and output rule set to be treated as certifiable software artifacts under DO-178C while the underlying ML ensemble is governed by a separate, less burdensome assurance regime applicable to advisory rather than flight-critical functions.

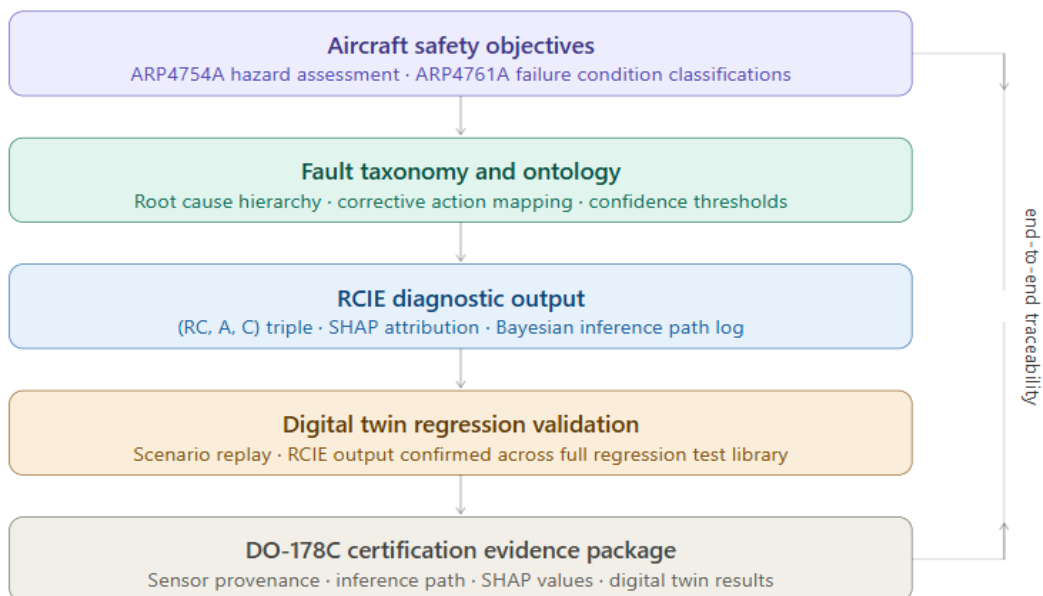


Figure 7: DMF Certification Evidence Flow

6.4

Limitations

The current DMF evaluation and architectural specification operate within boundaries that define the scope of the claims made in this article. Several factors remain outside the current scope and represent directions for future work rather than resolved design questions.

The evaluation is retrospective, drawing on historical maintenance records rather than prospective live-fleet operational data. The improvement magnitudes reported in Table 5 reflect simulation performance against verified historical ground truth and do not constitute evidence of equivalent performance in a live operational deployment where observation set completeness, data latency, and technician workflow integration introduce additional variables.

The fault taxonomy and RCIE ontology used in the evaluation cover six avionics fault categories representative of data link, hydraulic, engine thermal, landing gear, power supply, and navigation database failure modes. Coverage of the full operational fault space of a modern commercial transport avionics suite, including integrated modular avionics platform-specific failure modes and software-defined radio system faults, requires ontology extension work not included in the current implementation.

The closed-loop learning evaluation assumed a single-operator fleet with consistent maintenance documentation practices. Cross-operator federated learning, in which multiple airlines contribute to shared RCIE model improvements without exposing proprietary fleet data, is described architecturally in Section 5.6 but has not been empirically validated.

The current AISIP safety set definition assumes convex fault output spaces. Non-convex operational regions arising from multimodal fault signatures that span discontinuous regions of the feature space are not handled by the current projection formalism and represent a known limitation for fault categories with complex symptom interactions.

Adversarial or systematically corrupted input data scenarios, including sensor spoofing and bus protocol injection attacks, are outside the current threat model and would require integration with a separate cybersecurity monitoring layer before the

DMF could be deployed on platforms with elevated cyber threat exposure.

Conclusion

This article has presented the Deterministic Maintenance Framework as a formally specified, certifiably structured approach to avionics fault diagnosis that resolves the fundamental tension between the diagnostic sensitivity of AI/ML methods and the output determinism required for safety-critical aerospace certification. By defining maintenance diagnosis as the function $f(O) = (RC, A, C)$ and enforcing injectivity through the RCIE, the DMF guarantees that the same fault always yields the same root cause, the same corrective action, and the same auditable explanation, regardless of technician identity, geographic location, or model version history. The ten-layer architecture distributes this guarantee across a pipeline that is modularly decomposed for certification evidence generation, incrementally deployable for phased fleet introduction, and continuously improved through a constrained closed-loop learning mechanism validated by digital twin simulation before any update reaches the operational system. The quantitative results reported in Section 6 demonstrate that the DMF achieves operationally meaningful improvements across all measured dimensions: MTTR reduced by 57 percent, NFF rate reduced by 74 percent, first-time fix rate improved by 33 percentage points, and diagnostic consistency elevated above 99 percent from a baseline of 61 percent. These improvements translate directly into reduced aircraft on ground time, lower maintenance labor costs, and a more complete airworthiness documentation record. The framework's architectural alignment with DO-178C traceability principles and its structured response to the certification guidance articulated in the EASA AI Roadmap 2.0 and the FAA AI Safety Assurance Roadmap position it as a viable pathway for operators seeking to deploy AI assisted maintenance diagnostics under current regulatory frameworks without waiting for new certification standards to mature. Future research directions include developing standardized digital twin fidelity metrics for cross-operator benchmarking of DMF diagnostic rule validation quality, extending the RCIE ontology to cover emerging avionics architectures like

integrated modular avionics platforms and software-defined radio systems, and investigating federated learning mechanisms that let multiple airlines contribute to shared RCIE model improvements without exposing proprietary fleet data or creating cross-operator certification dependencies. The deterministic maintenance paradigm introduced here is not a final architecture but a foundational framework whose principles, guaranteed output reproducibility, causal explainability, and certification-aligned modularity are applicable wherever AI assisted diagnosis must meet a regulatory standard of accountability that probabilistic methods alone cannot satisfy.

References

- [1] R. K. Mobley, *An Introduction to Predictive Maintenance*. Elsevier, 2002. Available: <https://books.google.com/books?id=SjqXzxpAzSQ> C
- [2] 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. Available: <https://doi.org/10.1016/j.ymsp.2005.09.012>
- [3] J. Lee, B. Bagheri, and H.-A. Kao, "A cyber-physical systems architecture for industry 4.0-based manufacturing systems," *Manufacturing Letters*, vol. 3, pp. 18-23, 2015. Available: <https://doi.org/10.1016/j.mfglet.2014.12.001>
- [4] K. P. Murphy, *Machine Learning: A Probabilistic Perspective*. MIT Press, 2012. Available: <https://www.cs.ubc.ca/~murphyk/MLbook/>
- [5] J. Pearl, *Probabilistic Reasoning in Intelligent Systems*. Elsevier, 2014. Available: <https://dl.acm.org/doi/book/10.5555/534975>
- [6] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016.
- [7] F. Tao et al., "Data-driven smart manufacturing," *Journal of Manufacturing Systems*, vol. 48, pp. 157-169, 2018. Available: <https://doi.org/10.1016/j.jmsy.2018.01.006>
- [8] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," *Advances in Neural Information Processing Systems*, vol. 30, 2017. Available: <https://proceedings.neurips.cc/paper/2017/file/8a20a8621978632d76c43dfd28b67767-Paper.pdf>
- [9] M. Chen, R. Qu, and W. Fang, "Case-based reasoning system for fault diagnosis of aero-engines," *Expert Systems with Applications*, vol. 202, p. 117350, 2022. Available: <https://doi.org/10.1016/j.eswa.2022.117350>
- [10] Z. She et al., "Ontology-based root cause fault diagnostic reasoning for airborne systems," *Aerospace*, vol. 13, no. 2, p. 167, 2026. Available: <https://doi.org/10.3390/aerospace13020167>
- [11] S. Sutthithatip, S. Perinpanayagam, and S. Aslam, "(Explainable) artificial intelligence in aerospace safety critical systems," in *Proc. IEEE Aerospace Conference*, 2022, pp. 1-12. Available: <https://doi.org/10.1109/AERO53065.2022.9843612>
- [12] EASA, *Artificial Intelligence Roadmap 2.0*, May 2023. Available: <https://www.easa.europa.eu/en/document-library/general-publications/easa-artificial-intelligence-roadmap-20>
- [13] FAA, *Roadmap for Artificial Intelligence Safety Assurance, Version I*, Jul. 2024. Available: https://www.faa.gov/aircraft/air_cert/step/roadmap_for_AI_safety_assurance
- [14] W. K. Youn et al., "Software certification of safety critical avionic systems: DO-178C and its impacts," *IEEE Aerospace and Electronic Systems Magazine*, vol. 30, no. 4, pp. 4-13, 2015. Available: <https://doi.org/10.1109/MAES.2014.140109>
- [15] M. J. Scott et al., "A systematic literature review of predictive maintenance for defence fixed-wing aircraft sustainment and operations," *Sensors*, vol. 22, no. 18, p. 7070, 2022. Available: <https://doi.org/10.3390/s22187070>
- [16] J. Lee and M. Mitici, "Deep reinforcement learning for predictive aircraft maintenance using probabilistic remaining-useful life prognostics," *Reliability Engineering and System Safety*, vol. 230, p. 108908, 2023. Available: <https://doi.org/10.1016/j.res.2022.108908>
- [17] L. Li et al., "Digital twin in aerospace industry: A gentle introduction," *IEEE Access*, vol. 10, pp. 9543-9562, 2021. Available: <https://doi.org/10.1109/ACCESS.2021.3136458>
- [18] G. M. Bisanti et al., "Digital twins for aircraft maintenance and operation," *Internet of Things*, vol.

24, p. 100991, 2023. Available: <https://doi.org/10.1016/j.ijot.2023.100991>

[19] I. Kabashkin and V. Perekrestov, "Ecosystem of aviation maintenance," *Applied Sciences*, vol. 14, no. 11, p. 4394, 2024. Available: <https://doi.org/10.3390/app14114394>

[20] F. S. Cetin et al., "Data-driven predictive maintenance for aircraft components through sparse event logs," *Aerospace*, vol. 13, no. 1, p. 110, 2026. Available: <https://doi.org/10.3390/aerospace13010110>

[21] J. Li, S. King, and I. Jennions, "Intelligent fault diagnosis of an aircraft fuel system using machine learning," *Machines*, vol. 11, no. 4, p. 481, 2023. Available: <https://doi.org/10.3390/machines11040481>

[22] P. Andrade et al., "Aircraft maintenance check scheduling using reinforcement learning," *Aerospace*, vol. 8, no. 4, p. 113, 2021. Available: <https://www.mdpi.com/2226-4310/8/4/113>

[23] C. Silva et al., "Adaptive reinforcement learning for task scheduling in aircraft maintenance," *Scientific Reports*, vol. 13, no. 1, p. 16605, 2023. Available: <https://www.nature.com/articles/s41598-023-41169-3>

[24] A. Zaoui et al., "Impact of artificial intelligence on aeronautics: An industry-wide review," *Journal of Engineering and Technology Management*, vol. 71, p. 101800, 2024. Available: <https://doi.org/10.1016/j.jengtecman.2024.101800>

[25] M. Baptista et al., "Forecasting fault events for predictive maintenance using data-driven techniques and ARMA modeling," *Computers and Industrial Engineering*, vol. 115, pp. 41-53, 2018. Available: <https://doi.org/10.1016/j.cie.2017.10.033>

[26] R. Meissner et al., "Developing prescriptive maintenance strategies in the aviation industry based on a discrete-event simulation framework for post-prognostics decision making," *Reliability Engineering & System Safety*, vol. 214, pp. 107812, 2021. Available: <https://doi.org/10.1016/j.ress.2021.107812>