



AI-SAFE FlightNet: A Next-Generation Artificial Intelligence and Automation Framework for Predictive Flight Safety, Aircraft Health Intelligence, And Ground Operations Coordination

Prudvi Saisaran Ponduru¹, Pavani Priya Vyshnavi Nandanavanam², Sai Kesav Kumar Ponduru³
^{1,2,3}Researcher

ABSTRACT: Aviation safety improvement remains a continuous engineering challenge because modern aircraft are sensor-dense, software-intensive, and operationally interconnected, while health and operations data often remain fragmented across onboard, maintenance, dispatch, and ground systems. This fragmentation can delay weak-signal recognition and coordinated response. This paper proposes AI-SAFE FlightNet, a Safety-Aware, Intelligent, Federated, Explainable Flight Automation Network that integrates aircraft data intelligence, onboard edge AI, subsystem digital twins, phase-aware predictive analytics, evidence-grounded coordination, ground operations support, federated fleet learning, and governance by design. The framework supports pilots, engineers, dispatchers, and maintenance controllers without replacing certified human authority or certified aviation systems. A formal model combines sensor, flight-phase, environmental, and maintenance-history variables to estimate risk trajectories, prioritize alerts, and generate auditable response packages. The proposed evaluation uses public engine-degradation and trajectory datasets, synthetic subsystem telemetry, simulated maintenance records, and digital-twin experiments. Expected contributions are earlier warning, improved maintenance readiness, faster aircraft-to-ground coordination, privacy-preserving fleet learning, explainable recommendations, and measurable operational risk reduction. The objective is improved resilience and decision support, not zero risk or unrestricted autonomous control.

KEYWORDS: Artificial Intelligence; Aviation Safety; Predictive Maintenance; Digital Twin; Federated Learning; Explainable AI

I. INTRODUCTION

Aviation is a safety-critical socio-technical system in which aircraft design, operations control, maintenance, dispatch, airport readiness, crew decision-making, and regulatory oversight interact under uncertainty. Continuous safety improvement is therefore not a one-time technological milestone; it is a permanent requirement of system engineering, data governance, and operational discipline. Safety Management Systems and State Safety Programmes emphasize proactive hazard identification, risk assessment, monitoring, and continuous improvement [1], [2].

Modern aircraft systems are complex because many subsystems operate concurrently across different physical, computational, and organizational layers. Propulsion, flight controls, hydraulics, electrical distribution, fuel, avionics, cabin systems, environmental control, navigation, communication, maintenance records, and ground operations all generate safety-relevant information. Each signal may be individually normal while a multivariate pattern may indicate early degradation when interpreted with flight phase, environmental context, maintenance history, and operating condition.

Aircraft and airline operations now produce large volumes of sensor, avionics, maintenance, weather, operational, and

performance data. However, monitoring frequently remains distributed across onboard health systems, flight data

monitoring systems, maintenance information systems, operations control dashboards, dispatch systems, and airport teams. The resulting fragmentation can delay coordinated interpretation and can make it difficult to identify weak signals before operational risk increases.

Artificial intelligence can support improvement by detecting weak signals, forecasting degradation, estimating risk trajectories, ranking alerts, retrieving approved evidence, and generating structured coordination messages. The use of AI must be bounded by safety assurance, explainability, auditability, cybersecurity, and human authority. In this paper, AI is treated as a decision-support and coordination technology rather than a replacement for certified pilots, engineers, dispatchers, or certified flight-critical systems [3]-[7].

The research gap addressed here is the absence of an integrated, certifiable, human-supervised AI framework that connects aircraft health intelligence, predictive safety analytics, aircraft digital twins, ground notification, fleet-wide learning, and governance into one safety-aware automation network. The main contributions are: a novel architecture called AI-SAFE FlightNet; a formal system model for phase-aware risk estimation and ground response;

“AI-SAFE FlightNet: A Next-Generation Artificial Intelligence and Automation Framework for Predictive Flight Safety, Aircraft Health Intelligence, And Ground Operations Coordination”

algorithms for weak-signal fusion, digital twin simulation, federated fleet learning, and human-in-the-loop approval; safe automation levels; and an experimental design with reproducible benchmark data.

II. BACKGROUND AND RESEARCH FOUNDATIONS

A. Historical Evolution of Flight Automation and Safety Technology

The development of flight automation and aviation safety technology may be understood as a gradual transition from

mechanical robustness toward digital sensing, integrated monitoring, predictive reasoning, and safety-governed intelligence. The progression is described as a technology evolution rather than as a discussion of any event or operational occurrence.

TABLE I. EVOLUTION OF FLIGHT AUTOMATION AND SAFETY TECHNOLOGY

Phase	Mechanism	Safety Value	Remaining Limitation
Mechanical Control Systems	Mechanical linkages, cables, pulleys, trim devices	Direct control authority and basic controllability	Minimal automated monitoring
Hydraulic Systems	Power-assisted actuation and redundancy	Greater control force management and stability support	Fault context remained subsystem-local
Radio Navigation	Navigation beacons and radio-based guidance	Improved route guidance and position awareness	Limited integration with health intelligence
Autopilot Systems	Automatic attitude, heading, altitude, and path control	Reduced workload and improved flight path precision	Automation supervision still required
Jet-Age Monitoring Systems	Engine instruments, warning annunciation, and maintenance checks	Better propulsion and system awareness	Mostly threshold-based alerting
Flight Recorders	Data Recorded flight parameters and operational data	Post-flight analysis and safety monitoring	Often delayed analysis
Weather Radar	Onboard weather sensing and avoidance support	Improved environmental awareness	Interpretation remained crew-centered
Ground Proximity Warning Systems	Terrain proximity alerting logic	Enhanced terrain awareness	Limited predictive subsystem integration
Traffic Collision Avoidance Systems	Traffic surveillance and resolution advisories	Improved traffic conflict awareness	Specific operational scope
Fly-By-Wire	Digital flight control and envelope protections	Software-mediated control protections	Certification complexity
Full Authority Digital Engine Control	Digital engine control and monitoring	Optimized propulsion management	Limited cross-fleet learning
Aircraft Health Monitoring	Subsystem condition monitoring	Earlier maintenance awareness	Fragmented across systems
Predictive Maintenance	Data-driven degradation forecasting	Improved maintenance planning	Data access and validation constraints
Digital Twins	Synchronized computational asset models	Risk trajectory simulation	Fidelity and latency trade-offs
AI-Enabled Intelligence	Edge AI, explainable AI, federated learning, workflow automation	Weak-signal detection and coordinated response	Assurance, cybersecurity, certification, and human factors

B. Research Problem

Modern aircraft and airline operations generate massive real-time data, but safety intelligence is often distributed

across onboard systems, maintenance systems, dispatch systems, airport teams, and airline operations centers. The missing capability is an integrated AI-driven framework that

can identify early degradation, forecast risk, explain possible system behavior, notify the right human teams, and recommend safe response options before operational risk increases. This problem is not merely a modeling problem; it

is an aircraft-ground coordination, governance, and human-supervised automation problem [20]-[24].

C. Research Questions

TABLE II. RESEARCH QUESTIONS

ID	Research Question
RQ1	How can AI identify early indicators of aircraft system degradation?
RQ2	How can edge AI support real-time onboard aircraft health monitoring?
RQ3	How can aircraft digital twins simulate possible system behavior before risk escalates?
RQ4	How can AI automate structured communication between aircraft systems and ground operations?
RQ5	How can fleet-wide federated learning improve safety intelligence while protecting airline data?
RQ6	How can AI recommendations remain explainable, auditable, and human-supervised?
RQ7	What new automation layer is required for future intelligent flight safety systems?

D. Literature Review

Aviation safety management literature emphasizes proactive hazard identification, risk mitigation, assurance, and organizational accountability. ICAO Annex 19 and the Safety Management Manual describe safety management as a continuous organizational process that must combine policy, risk management, assurance, and promotion [1], [2]. This establishes the governance rationale for data-driven safety intelligence.

Aircraft health monitoring and prognostics literature shows that aircraft maintenance is moving from isolated component inspection toward platform health management and predictive maintenance. Reviews of aircraft prognostic and health management identify opportunities in diagnosis, prognosis, health monitoring, remaining useful life estimation, and decision support, while also warning about data quality, rare failure labels, integration, and validation difficulty [20]-[24].

Flight data monitoring research shows that routine operational data can produce safety knowledge when analyzed systematically. Machine learning has been used for anomaly detection in flight data monitoring, but the challenge remains to convert anomalies into explainable, role-specific, and operationally useful decisions [25], [26].

Fault detection and isolation research provides the diagnostic foundation for detecting abnormal behavior from residuals, models, data histories, and multivariate sensor patterns. Classical fault diagnosis emphasizes residual generation and structured isolation, while current data-driven methods include deep sequence learning and graph-based reasoning for complex systems [27]-[32].

Digital twin literature in aerospace argues that a useful twin is a synchronized computational representation that combines physics, operating state, and data history to support prediction and decision-making. Digital twins can support

risk trajectory forecasting, what-if simulation, subsystem interaction analysis, and maintenance planning [17]-[19].

AI time-series anomaly detection provides tools for learning degradation from multivariate sensor sequences. Random forests, gradient boosting, recurrent neural networks, temporal convolutional networks, transformers, autoencoders, graph neural networks, and Bayesian networks each provide different trade-offs in interpretability, temporal memory, computational cost, and uncertainty handling [33]-[41].

Edge AI is relevant because onboard and near-aircraft computation can reduce latency, preserve bandwidth, and maintain limited inference capability during connectivity disruption. Federated learning is relevant because fleet-wide learning can improve model generalization while reducing the need to centralize proprietary raw data [42]-[44].

Explainable AI and human factors research are essential because aviation decisions require appropriate reliance, traceability, and human understanding. LIME, SHAP, and broader XAI research support explanation, while automation research shows that poor trust calibration, alert fatigue, and out-of-the-loop performance can reduce safety benefit if automation is not carefully designed [45]-[50].

Safety-critical AI assurance and certification research indicates that performance metrics alone are insufficient. FAA and EASA guidance emphasizes AI safety assurance, human-centric design, defined operating domains, traceable requirements, and incremental adoption for machine learning applications [3]-[6], [51]. Cybersecurity literature further implies that connected aircraft-ground systems require zero-trust principles, secure communication, model integrity, and auditability [7]-[9].

III. AI-SAFE FLIGHTNET ARCHITECTURE

A. Proposed Framework

AI-SAFE FlightNet is proposed as a Safety-Aware, Intelligent, Federated, Explainable Flight Automation Network. The framework is designed as a multi-layer architecture that connects aircraft sensing, onboard inference,

digital twin simulation, predictive safety analytics, agentic coordination, ground operations, fleet-wide learning, and governance. The system supports human decision makers and certified aviation systems; it does not replace certified human authority or certified flight-critical functions.

TABLE III. AI-SAFE FLIGHTNET MULTI-LAYER ARCHITECTURE

Layer	Core Components	Safety Function
Aircraft Data Intelligence Layer	Engine sensor data; flight control data; hydraulic data; electrical data; fuel data; cabin data; avionics data; environmental data; maintenance history data	Creates a complete context-aware state picture
Onboard Edge AI Layer	Real-time preprocessing; sensor fusion; lightweight anomaly detection; flight-phase-aware risk scoring; local inference; certified alert filtering	Provides low-latency detection and first-pass triage
Aircraft Digital Twin Layer	Engine twin; flight control twin; electrical twin; fuel twin; environmental performance model; failure propagation simulation; risk trajectory forecasting	Simulates future behavior and escalation paths
Predictive Safety Intelligence Layer	Time-series forecasting; weak-signal detection; fault pattern recognition; risk ranking; maintenance prediction; remaining useful life estimation; alert severity classification	Transforms telemetry into predictive safety insight
Agentic AI Coordination Layer	AI safety assistant; reasoning engine; evidence retrieval from approved manuals and databases; recommendation generator; human approval gateway; structured communication generator; ground workflow automation	Converts risk insight into explainable human-supervised action proposals
Ground Operations Intelligence Layer	Operations control dashboard; maintenance control alerting; engineering notification; dispatcher decision support; airport readiness coordination; spare parts and inspection planning; post-flight technical reporting	Accelerates coordinated ground response
Fleet-Wide Learning Layer	Federated learning; privacy-preserving model updates; cross-fleet anomaly learning; model drift detection; safety trend analytics; continuous improvement loop	Learns subtle fleet-wide degradation patterns
Governance And Certification Layer	Human authority; explainability; audit logging; safety policy enforcement; cybersecurity controls; model validation; fail-safe boundaries; certification readiness	Maintains assurance, accountability, and deployment boundaries

B. Architecture Diagram




Fig. 1. AI-SAFE FlightNet Architecture Diagram

C. AI Reasoning Loop Diagram




Fig. 2. AI Reasoning Loop Diagram

D. Cockpit-to-Ground Automation Workflow




Fig. 3. Cockpit-To-Ground Automation Workflow

E. Novel Research Contribution

The core novelty is Federated Weak-Signal Safety Intelligence For Predictive Aircraft Risk Reduction. The proposed system does not only monitor one aircraft. It learns subtle degradation patterns across an entire fleet using privacy-preserving AI, edge intelligence, aircraft digital twins, and human-supervised automation. A single aircraft may not provide enough examples of rare incipient degradation, but a fleet-wide federated learning system can learn shared representation patterns without requiring raw proprietary data to be centralized [42]-[44].

The framework is novel because its output is not merely an anomaly score. Its output is an auditable safety-intelligence package that includes phase-aware risk, supporting evidence, digital twin forecast, role-specific notification, suggested inspection, confidence level, and human validation state.

IV. FORMAL METHODS AND AUTOMATION DESIGN

A. Formal System Model

Let the latent aircraft state vector at time t be represented as x_t in $\mathbb{R}^n$, where the state spans propulsion, flight control, electrical, hydraulic, fuel, cabin, avionics, and environmental subsystems.

Let the sensor observation vector be y_t in $\mathbb{R}^m$. Let the flight phase variable be ϕ_t in Φ , where Φ includes taxi, takeoff, climb, cruise, descent, approach, landing, and turnaround. Let the environment vector be e_t in $\mathbb{R}^q$ and the maintenance history vector be h_t in $\mathbb{R}^p$.

The predictive risk score is defined as shown in Equation 1.

$$r_t = f_{\theta}(y_{1:t}, \phi_t, e_t, h_t) \quad (1)$$

The estimated near-horizon failure or unsafe-degradation probability is shown in Equation 2.

$$p_hat_t = g_theta(y_1:t, x_t, phi_t, e_t, h_t) \quad (2)$$

A phase-dependent alert threshold τ_{ϕ_t} is used because risk meaning changes across flight phases. An alert is issued when the condition in Equation 3 holds.

$$I(r_t \geq \tau_{\phi_t}) = 1 \quad (3)$$

The ground response function Γ maps risk, probability, flight phase, explanation, and safety policy into a recommended action package.

$$a_t = \Gamma(r_t, p_hat_t, \phi_t, explain_t, policy) \quad (4)$$

The human validation function H maps the recommended action into approve, modify, reject, or request-more-evidence decisions.

$$u_t = H(a_t, context_t, authority_t) \quad (5)$$

The optimization objective minimizes missed detection cost, false alarm cost, latency, drift, and weak explanation quality while rewarding useful early warning. It is not an objective of zero risk.

$$\min_{\theta} E[\lambda_{FN} CFN + \lambda_{FP} CFP + \lambda_L L_{latency} + \lambda_D L_{drift} - \lambda_E U_{early-warning}] \quad (6)$$

For automation levels 0 through 4, the safety constraint is that unsafe operational recommendations must not be executed without certified human validation.

$$Pr(unsafe\ recommendation\ without\ human\ validation) = 0 \quad (7)$$

B. AI and Automation Models

TABLE IV. AI AND AUTOMATION MODELS

Model Family	Use In AI-SAFE FlightNet	Strength	Caution
Random Forest	Tabular degradation classification and feature ranking	Robust and interpretable feature importance	Limited temporal memory
Gradient Boosting	Fault severity and maintenance action ranking	Strong structured-data performance	Requires calibration
Long Short-Term Memory	Sequential degradation tracking	Long-range temporal state	Opaque and compute-intensive
Gated Recurrent Unit	Edge-friendly temporal inference	Lower complexity than LSTM	May capture less memory
Temporal Convolutional Network	Fast multivariate sequence modeling	Parallelizable and stable receptive field	Window design matters
Time-Series Transformer	Long-horizon dependency learning	Captures cross-channel attention	Requires more data and compute
Autoencoder	Unsupervised anomaly detection	Useful with sparse labels	Reconstruction error can be miscalibrated
Graph Neural Network	Inter-subsystem dependency reasoning	Captures relational propagation	Graph design is difficult
Bayesian Network	Uncertainty-aware diagnosis	Transparent probability propagation	Structure learning is demanding
Reinforcement Learning	Maintenance and ground operations optimization	Optimizes sequential decisions	Requires strict safety constraints
LLM Agent Retrieval	With Approved documentation retrieval and structured communication	Improves usability and drafting	Must be bounded and human-approved
Digital Twin Simulation	Failure propagation and scenario forecasting	Physics-informed anticipation	Fidelity and latency trade-off

C. Algorithms

1) Real-Time Aircraft Anomaly Detection:

Input: sensor stream y_t , phase ϕ_t , environment e_t

Output: anomaly score a_t and explanation z_t

1: buffer $\leftarrow$ update_window(y_t, ϕ_t, e_t)

2: $y_clean \leftarrow$ preprocess(buffer)

3: $y_fused \leftarrow$ sensor_fusion(y_clean, ϕ_t, e_t)

4: $a_t \leftarrow$ anomaly_model.infer(y_fused)

5: $z_t \leftarrow$ local_explainer(a_t, y_fused)

6: if $a_t > \text{threshold}(\phi_t)$ then emit anomaly_event(a_t, z_t, ϕ_t)

7: return a_t, z_t

2) Weak-Signal Fusion:

Input: channel scores $s_1 \dots s_k$, subsystem graph G , history h_t

Output: fused risk r_t

1: contextualize every channel score by flight phase and history

2: estimate persistence across time windows

```

3: propagate weak dependencies across subsystem graph
G
4: compute cross-channel consistency weights
5: r_t <- weighted_aggregate(channel scores, persistence,
graph effects)
6: return calibrated(r_t)
3) Predictive Maintenance Forecasting:
Input: multivariate sequence Y, maintenance history H
Output: remaining useful life R_hat and maintenance class
m_hat
1: features <- temporal_encoder(Y, H)
2: R_hat <- RUL_model(features)
3: m_hat <- maintenance_classifier(features, R_hat)
4: uncertainty <- estimate_uncertainty(features)
5: return R_hat, m_hat, uncertainty
4) Digital Twin Risk Simulation:
Input: state estimate x_t, degradation hypotheses Omega
Output: ranked risk trajectories T
1: twin_state <- synchronize(x_t)
2: for each hypothesis omega in Omega do
3:   simulate subsystem evolution over horizon H
4:   estimate propagation to coupled systems
5:   compute risk curve rho_omega(1:H)
6: rank trajectories by severity, likelihood, and
intervention time
7: return T
5) Federated Fleet Learning:
Input: local aircraft datasets D_i, global model theta
Output: validated global model theta_star
1: broadcast theta to participating fleet nodes
2: for each node i in parallel do
3:   theta_i <- local_train(theta, D_i)
4:   delta_i <- privacy_preserve(theta_i - theta)
5: discard suspicious or poisoned updates
6: theta_star <- secure_aggregate(delta_1...delta_n)
7: perform drift, robustness, and safety validation
8: publish validated advisory model
6) AI-Based Ground Staff Notification:
Input: risk package P, evidence E, policy rules R
Output: role-specific notification set N
1: infer impacted roles from affected system and severity

```

```

2: retrieve approved maintenance and operations
references
3: generate structured notification fields
4: attach explanation summary and audit trace
5: validate against policy rules R
6: route to appropriate human queues
7: return N
7) Alert Prioritization:
Input: active alert set A
Output: priority queue Q
1: for each alert a in A do
2:   score(a) = severity × confidence × phase_criticality
× persistence
3:   reduce score when explanation quality is low
4:   adjust score using maintenance history and
operational context
5: Q <- sort_descending(A, score)
6: return Q
8) Human-in-the-Loop Approval:
Input: recommendation r, evidence pack e, user authority
u
Output: signed decision d
1: display recommendation, confidence, explanation, and
alternatives
2: require approve, modify, reject, or request-more-
evidence decision
3: capture rationale, timestamp, and authority role
4: log all interaction to immutable audit store
5: return signed human decision d
9) Continuous Post-Flight Learning:
Input: completed flight data F, maintenance outcomes M,
human decisions U
Output: updated learning artifacts
1: align flight, maintenance, and decision logs
2: label confirmed anomalies, false alarms, and missed
detections
3: update performance dashboard and drift monitors
4: curate training examples for local and federated
retraining
5: refresh explanation libraries and alert rules
6: archive artifacts for audit and assurance

```

D. Safe Automation Levels

TABLE V. SAFE AUTOMATION LEVELS

Level	Definition	Typical Function	Human Authority
Level 0	Data monitoring	Recording, dashboarding, telemetry access	Full human interpretation
Level 1	Advisory alerts	Threshold alerts, anomaly flags, trend notices	Human reviews and decides
Level 2	Explainable recommendations	Ranked hypotheses and recommended inspections	Human validates recommendation
Level 3	Ground workflow automation	Draft messages, inspection tasks, spare readiness	Human supervises and approves
Level 4	Certified safety assistance	Bounded support integrated with approved concepts	Human remains final authority

Level	Definition	Typical Function	Human Authority
Level 5	Autonomous intervention only if formally certified and regulator-approved	Narrow autonomous action in certified design domain	Oversight framework remains mandatory

AI-SAFE FlightNet is primarily intended for Levels 1 through 4. Level 5 is included only as a theoretical future state requiring explicit regulator approval, formal certification, and bounded operating conditions. The framework does not claim autonomous replacement of pilots, engineers, dispatchers, or certified systems.

E. Ground Staff Notification System

Ground coordination is a first-class safety function. AI-SAFE FlightNet generates structured information for maintenance control, engineering teams, dispatchers, airline operations control, airport readiness teams, inspection crews, spare parts teams, and safety management teams. Notifications are role-specific, evidence-grounded, auditable, and subject to human validation.

TABLE VI. SAMPLE GROUND STAFF NOTIFICATION FIELDS

Field	Description
Aircraft Identifier	Internal aircraft or fleet identifier
Flight Phase	Taxi, takeoff, climb, cruise, descent, approach, landing, or turnaround
System Affected	Engine, hydraulic, electrical, flight control, fuel, cabin, avionics, or environmental
Risk Severity	Low, moderate, high, or critical
Confidence Score	Calibrated confidence with uncertainty flag
Supporting Sensor Evidence	Key channels, trends, residuals, and model explanations
Recommended Inspection	Suggested bounded follow-up inspection or check
Required Specialist Team	Relevant engineering or maintenance specialty
Estimated Urgency	Immediate, before next departure, or within maintenance window
Suggested Ground Action	Prepare tools, spares, inspection crew, dispatch consultation, or safety review
Audit Trace ID	Immutable trace identifier for the decision and communication chain

V. EXPERIMENTAL DESIGN AND EVALUATION

A. Experimental Design

A reproducible evaluation should combine public benchmark data, open trajectory data, synthetic subsystem telemetry, simulated maintenance logs, and digital twin

simulation. This mixed setup addresses a key limitation in aviation AI research: public access is stronger for propulsion and trajectories than for complete integrated aircraft health systems [10]-[16].

TABLE VII. EXPERIMENTAL DATA AND SIMULATION SOURCES

Source	Use In Evaluation
NASA C-MAPSS Turbofan Engine Degradation Dataset	Benchmark remaining useful life estimation and degradation forecasting
N-CMAPSS Dataset	Run-to-failure engine degradation under more realistic operating conditions
OpenSky Trajectory Data	Trajectory context, flight phase features, and operational context
ADS-B Trajectory Data	Trajectory reconstruction, phase labeling, and communication latency context
Synthetic Aircraft System Telemetry	Hydraulic, electrical, fuel, cabin, and avionics anomaly scenarios
Simulated Maintenance Logs	Inspection history, deferred actions, component age, and work package outcomes
Digital Twin Simulation Environment	Subsystem degradation simulation, risk trajectory forecasting, and intervention timing

B. BASELINES

TABLE VIII. BASELINE METHODS

Baseline	Description		Limitation Relative To AI-SAFE FlightNet
Rule-Based Monitoring	Threshold	Static limits on selected sensor channels	Weak multivariate signal sensitivity
Traditional Maintenance	Predictive	Component-focused RUL models	Limited cross-system and ground coordination
Single-Model Detection	Anomaly	One model used across contexts	Weak adaptability to phase and subsystem heterogeneity
Cloud-Only AI	Centralized inference and analytics		Latency and connectivity dependence
Edge-Only AI	Local aircraft inference only		Limited fleet learning and long-horizon simulation
Digital Twin Only	Simulation without federated learning or workflow automation		Limited adaptation from fleet history
Federated Learning Only	Privacy-preserving model updates		Weak scenario simulation and action orchestration
AI-SAFE FlightNet	Integrated edge, twin, predictive, agentic, ground, fleet, and governance layers		Higher implementation and assurance complexity

C. Evaluation Metrics

TABLE IX. EVALUATION METRICS

Category	Metrics
Predictive Performance	Accuracy, precision, recall, F1-score, AUROC
Error Management	False positive rate and false negative rate
Timeliness	Mean time to detect, time to ground notification, communication latency
Operational Utility	Alert usefulness score and maintenance action readiness time
Explainability And Trust	Explainability score, human agreement rate, and override rate
Learning Stability	Model drift rate and federated convergence stability
Cybersecurity And Assurance	Cybersecurity risk score, audit completeness, and safety case completeness

VI. RESULTS AND DISCUSSION

A. Experimental Results

This paper proposes a framework and experimental design. No experimental results are invented. The following tables

are placeholders to be filled only after reproducible experiments are conducted.

TABLE X. PREDICTIVE PERFORMANCE RESULTS PLACEHOLDER

Model Or Baseline		Accuracy	Precision	Recall	F1-Score	AUROC
Rule-Based Monitoring	Threshold	TBD	TBD	TBD	TBD	TBD
Traditional Maintenance	Predictive	TBD	TBD	TBD	TBD	TBD
AI-SAFE FlightNet		TBD	TBD	TBD	TBD	TBD

TABLE XI. OPERATIONAL COORDINATION RESULTS PLACEHOLDER

Model Or Baseline	Mean Time To Detect	Time To Ground Alert Notification	Usefulness Score	Maintenance Readiness Time
Cloud-Only AI	TBD	TBD	TBD	TBD

Model Or Baseline	Mean Time To Detect	Time To Notification	Ground Alert Score	Usefulness	Maintenance Readiness Time	Action
Edge-Only AI	TBD	TBD	TBD		TBD	
AI-SAFE FlightNet	TBD	TBD	TBD		TBD	

B. Ablation Study

TABLE XII. ABLATION STUDY DESIGN

Ablation Condition	Question Tested	Primary Outcome
Without Edge AI	Does local inference reduce latency and nuisance alerts?	Detection latency and communication dependence
Without Digital Twin	Does simulation improve anticipation?	Early-warning horizon and scenario quality
Without Federated Learning	Does fleet learning improve weak-signal capture?	Generalization to rare patterns
Without Agentic AI Coordinator	Does coordination reasoning improve actionability?	Alert usefulness and human workload
Without Ground Notification Automation	Does structured automation reduce coordination delay?	Time to notification and readiness time
Without Explainability Layer	Do explanations improve human acceptance?	Human agreement and trust calibration
Rule-Based Alerting Versus AI Risk Scoring	Are context-aware scores superior to static thresholds?	Precision and recall trade-off
Single Aircraft Learning Versus Fleet Learning	Does federated weak-signal learning help?	Rare-pattern robustness

C. Safety, Certification, and Governance

AI-SAFE FlightNet must be deployed through incremental safety assurance. Early functions should remain outside autonomous command of flight-critical actuation and should be limited to monitoring, advisory, explanation, and structured ground workflow support. FAA and EASA AI guidance indicates that AI assurance requires bounded applications, traceability, data management, validation, and human-centric deployment [3]-[6].

Human authority must be explicit. Pilot decision support, engineer decision support, dispatcher decision support, and maintenance control support should be separated from certified authority. The system may rank evidence and draft recommendations, but the acceptance of risk-bearing decisions must remain with certified personnel unless a future function is formally certified and regulator-approved.

Governance should include model inventory, data lineage, operating domain definition, validation evidence, drift monitoring, explanation policy, cybersecurity controls, audit logs, fail-safe boundaries, override capture, and a safety case. False alarm management is essential because excessive nuisance alerts can reduce trust and increase workload.

D. Cybersecurity and Data Protection

Connected aircraft-ground AI must be designed against sensor spoofing, data poisoning, model poisoning, unauthorized access, ground system compromise, and tampered federated updates. Secure communication, device identity, encryption, access control, and continuous

verification are required. Zero-trust architecture and AI risk management principles provide a useful foundation for this design [7]-[9].

Federated learning protects data sovereignty by sharing model updates rather than raw data, but it does not automatically prevent poisoning or drift. AI-SAFE FlightNet therefore requires secure aggregation, anomaly screening of model updates, drift validation, signed model release, and rollback mechanisms before a model can be admitted into operational advisory workflows.

E. Human Factors

Human factors determine whether the framework improves safety or creates new workload. Alert fatigue, trust miscalibration, automation bias, out-of-the-loop performance, and supervisory overload are credible hazards in any advanced automation system [48]-[50].

Cockpit advisories must be concise, phase-aware, prioritized, and explainable. Ground staff decision support must be evidence-rich, role-specific, and auditable. Dispatcher situational awareness requires clear operational alternatives and uncertainty representation. Maintenance engineer usability requires diagnostic evidence, trend history, likely affected line-replaceable units, and suggested inspection logic.

The system should reduce workload by removing clerical burden, not by hiding uncertainty. Human-machine teaming is improved when the model shows why a risk is ranked, how

confident it is, what evidence supports the claim, what alternatives exist, and what human action is required.

F. Comparative Tables

TABLE XIII. TRADITIONAL MONITORING VERSUS AI-SAFE FLIGHTNET

Dimension	Traditional Monitoring	AI-SAFE FlightNet
Signal Processing	Channel-by-channel thresholding	Multivariate, phase-aware, fused weak-signal intelligence
Timing	Reactive or delayed	Predictive and proactive
Explainability	Threshold exceedance explanation	Feature attribution, evidence retrieval, and uncertainty
Coordination	Manual escalation	Structured role-specific automation
Learning Scope	Asset-local or program-local	Fleet-wide federated learning
Governance	Procedure-based	Procedure, model, audit, and cyber governance

TABLE XIV. CLOUD AI VERSUS EDGE AI VERSUS HYBRID AI

Dimension	Cloud AI	Edge AI	Hybrid AI-SAFE FlightNet
Latency	Higher	Lower	Low for detection and managed for learning
Connectivity Dependence	High	Low	Managed
Fleet Knowledge	Strong	Weak	Strong
Coordination	Moderate	Weak	Strong
Operational Robustness	Variable	Good locally	Balanced aircraft-ground resilience

TABLE XV. CURRENT LIMITATIONS VERSUS PROPOSED INNOVATION

Current Limitation	Proposed Innovation
Fragmented aircraft-ground visibility	Unified safety automation network
Threshold-driven alerts	Weak-signal multivariate intelligence
Late diagnosis	Predictive digital twin reasoning
Limited cross-fleet learning	Privacy-preserving federated learning
Opaque recommendations	Explainable and auditable recommendations
Manual coordination	Structured human-supervised notification automation

G. Findings

1) *Finding 1: Weak-signal, multivariate, context-aware safety intelligence is more suitable for modern aircraft health prediction than isolated threshold exceedance.*

Evidence from literature: PHM, predictive maintenance, and FDI reviews show a movement from threshold alarms toward multivariate trend modeling, residual reasoning, and sequence learning [20]-[32].

Technical explanation: Weak degradations may not be individually salient, but correlations across channels, phases, and histories can make them detectable.

Practical implication: Safety systems should fuse subsystem signals and operating context before triage.

Limitation: Weak-signal models may raise false alarms if context calibration is poor.

Future research direction: Develop calibrated phase-specific confidence estimation and uncertainty-aware alert suppression.

2) *Finding 2: Edge AI is important for timeliness, but edge-only intelligence is insufficient for fleet-scale safety learning.*

Evidence from literature: Edge inference supports latency, while federated learning supports distributed collaborative training under data sovereignty constraints [42]-[44].

Technical explanation: Immediate detection and long-horizon fleet learning operate on different time constants.

Practical implication: Hybrid aircraft-ground AI is preferable to purely edge or purely cloud architectures.

Limitation: Hybrid systems are harder to assure, synchronize, and secure.

Future research direction: Optimize latency, privacy, and assurance jointly.

3) *Finding 3: Digital twins are most valuable when used for risk trajectory forecasting, not only visualization.*

Evidence from literature: Aerospace digital twin literature emphasizes synchronized asset models and scenario projection [17]-[19].

Technical explanation: A predictive twin converts diagnosis into anticipation by simulating future subsystem evolution.

Practical implication: Twins should support alert ranking and maintenance planning.

Limitation: Twin fidelity and real-time computability must be balanced.

Future research direction: Develop adaptive-fidelity twins for safety operations.

4) *Finding 4: Explainability and human authority are enabling conditions for aviation AI adoption.*

Evidence from literature: XAI and human-automation research show that appropriate reliance depends on understandable rationale [45]-[50].

Technical explanation: Operators must evaluate what the system recommends, why it recommends it, and how certain it is.

Practical implication: Every advisory should include evidence, uncertainty, and provenance.

Limitation: Some explanation methods can be unstable or misleading.

Future research direction: Create aviation-specific explanation metrics tied to workload and decision quality.

5) *Finding 5: Coordination automation is an underdeveloped but high-impact safety lever.*

Evidence from literature: Safety management and maintenance literature show that value can be lost during organizational handoff [1], [2], [21]-[25].

Technical explanation: A correct diagnosis that arrives late, to the wrong team, or without structure has limited safety value.

Practical implication: Structured ground notifications should be treated as safety workflow outputs.

Limitation: Over-automation of workflow can create complacency.

Future research direction: Study notification routing, granularity, acknowledgment, and override design.

H. Discussion of Findings

AI can realistically improve anomaly detection, degradation forecasting, risk ranking, maintenance triage, documentation retrieval, structured communication, and post-flight learning today. These are domains where data are large,

decisions are repetitive but consequential, and human-supervised decision support can create measurable operational value.

What must remain human-supervised includes pilot operational decisions, engineering release decisions, dispatch decision finalization, maintenance control decisions, and any safety-critical command authority unless a function is separately certified. The framework therefore avoids unsafe automation dependence by staging AI outputs as advisory, explainable, auditable, and approval-gated.

Fleet-wide weak-signal detection is valuable because rare degradation precursors may be difficult to learn from one aircraft. Federated learning allows subtle patterns to be learned across a fleet while reducing raw data centralization. This approach can support continuous improvement while respecting data confidentiality.

A gradual deployment strategy is recommended: post-flight analytics first, near-real-time ground support second, bounded onboard advisory support third, and only then certified safety assistance. This sequence allows organizations to validate data quality, monitor false alarms, refine explanations, build trust, and construct safety cases before expanding automation scope.

I. Limitations

- Proprietary integrated flight and maintenance data are difficult to access for open research.
- Rare-event validation is difficult because safety-significant patterns occur infrequently.
- Regulatory constraints and certification complexity can slow deployment.
- Human factors uncertainty remains because trust and workload are interface-dependent.
- Dataset bias may occur when public datasets emphasize engines and trajectories more than full aircraft systems.
- Cybersecurity risks include spoofing, poisoning, unauthorized access, and model compromise.
- Model drift may occur as aircraft age, operations change, or sensor calibration shifts.
- Real-world validation, shadow-mode trials, and operational safety cases are required before deployment.

J. Ethics and Responsible AI

Responsible aviation AI begins with passenger safety, crew support, and preserved human authority. AI-SAFE FlightNet is therefore designed to support decision makers rather than replace them. It must not encourage overtrust, conceal uncertainty, or create pressure to accept model recommendations.

Crew privacy, data governance, airline confidentiality, and regulatory compliance must be maintained. The system should minimize sensitive data exposure, restrict access by role, preserve audit trails, and use privacy-preserving learning where feasible. Ethical use also requires that model performance, bias, drift, and cybersecurity posture be continuously monitored.

K. Future Research Agenda

TABLE XVI. FUTURE RESEARCH AGENDA

Area	Research Direction
Federated Aviation AI	Robust aggregation, poisoning resistance, cross-fleet domain adaptation
Edge AI For Onboard Health Monitoring	Certifiable low-latency models and bounded compute
Digital Twins For Aircraft Systems	Adaptive fidelity, online synchronization, uncertainty-aware simulation
Explainable AI For Pilots And Engineers	Role-specific explanation design and workload-aware explanation metrics
AI-Assisted Maintenance Planning	Joint RUL, spare parts, staffing, and inspection optimization
Ground Operations Automation	Human-supervised workflow orchestration and structured communication semantics
Certified LLM Agents	Retrieval-constrained agents using approved manuals and bounded outputs
Aviation AI Benchmarks	Open multimodal PHM and operations datasets with reproducibility protocols
Safety-Governed Aircraft Automation	Assurance cases, out-of-distribution monitoring, fail-safe policies, and staged certification
Human-AI Teaming	Trust calibration, alert design, supervisory handoff, and role-adaptive interfaces

VII. CONCLUSION

This paper proposed AI-SAFE FlightNet, a next-generation Artificial Intelligence and automation framework for predictive flight safety, aircraft health intelligence, and ground operations coordination. The framework integrates aircraft sensing, onboard edge AI, subsystem digital twins, predictive safety intelligence, agentic coordination, structured ground notification, federated fleet learning, and governance-by-design.

The central contribution is Federated Weak-Signal Safety Intelligence For Predictive Aircraft Risk Reduction. The framework is intended to enable earlier warning, predictive maintenance, ground staff coordination, fleet-wide learning, explainable recommendations, and human-supervised automation. It does not claim that risk can be reduced to zero. The intended goal is measurable risk reduction, faster coordination, better maintenance readiness, and improved operational resilience.

The path to aviation AI adoption should be incremental, certifiable, auditable, secure, and human-centered. Under these constraints, AI-SAFE FlightNet provides a research blueprint for the next generation of predictive aviation safety systems.

ACKNOWLEDGMENT

The authors acknowledge the broader aviation safety, aerospace systems, artificial intelligence, prognostics and health management, human factors, and safety assurance research communities whose published work informs the research direction of this paper.

CONFLICT OF INTEREST

The authors declare no known conflict of interest related to the proposed framework. No external sponsor influenced

the technical design, interpretation, or conclusions of this paper.

REFERENCES

1. International Civil Aviation Organization, Annex 19: Safety Management, ICAO, Montreal, Canada.
2. International Civil Aviation Organization, Safety Management Manual, Doc 9859, 4th ed., ICAO, Montreal, Canada, 2018.
3. Federal Aviation Administration, Roadmap for Artificial Intelligence Safety Assurance, Washington, DC, USA, July 2024.
4. Federal Aviation Administration, Technical Discipline: Artificial Intelligence - Machine Learning, Washington, DC, USA, 2026.
5. European Union Aviation Safety Agency, Artificial Intelligence Roadmap 2.0: A Human-Centric Approach To AI In Aviation, Cologne, Germany, 2023.
6. European Union Aviation Safety Agency, Artificial Intelligence Concept Paper Issue 2: Guidance For Level 1 And Level 2 Machine-Learning Applications, Cologne, Germany, 2024.
7. National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework, NIST AI 100-1, 2023.
8. RTCA, DO-326A Airworthiness Security Process Specification, Washington, DC, USA.
9. S. Rose, O. Borchert, S. Mitchell, S. Connelly, Zero Trust Architecture, NIST Special Publication 800-207, 2020.
10. National Aeronautics and Space Administration, CMAPSS Jet Engine Simulated Data, NASA Open Data Portal.

11. National Aeronautics and Space Administration, Prognostics Center Of Excellence Data Set Repository, NASA Ames Research Center.
12. A. Saxena, K. Goebel, D. Simon, N. Eklund, “Damage Propagation Modeling For Aircraft Engine Run-To-Failure Simulation”, Proceedings of the International Conference on Prognostics and Health Management, 2008.
13. M. Arias Chao, C. Kulkarni, K. Goebel, O. Fink, “Aircraft Engine Run-To-Failure Dataset Under Real Flight Conditions For Prognostics And Diagnostics”, Data, 2021, 6 (1), 5.
14. The OpenSky Network, OpenSky Network Data Access And Scientific Datasets.
15. M. Strohmeier, J. Schafer, V. Lenders, I. Martinovic, “Crowdsourced Air Traffic Data From The OpenSky Network 2019-20”, Earth System Science Data, 2021, 13, 357-366.
16. Federal Aviation Administration, Automatic Dependent Surveillance-Broadcast, Washington, DC, USA, 2025.
17. E.J. Tuegel, A.R. Ingraffea, T.G. Eason, S.M. Spottswood, “Reengineering Aircraft Structural Life Prediction Using A Digital Twin”, International Journal of Aerospace Engineering, 2011, 2011, 154798.
18. E.H. Glaessgen, D. Stargel, “The Digital Twin Paradigm For Future NASA And U.S. Air Force Vehicles”, Proceedings of the AIAA Structures, Structural Dynamics, and Materials Conference, 2012.
19. R. van Dinter, B. Tekinerdogan, C. Catal, “Predictive Maintenance Using Digital Twins: A Systematic Literature Review”, Information and Software Technology, 2022, 151, 107008.
20. A. Theissler, M. Perez-Velazquez, M. Kettelgerdes, G. Elger, “Predictive Maintenance Enabled By Machine Learning: Use Cases And Challenges”, Reliability Engineering And System Safety, 2021, 215, 107864.
21. J. Lee, M. Mitici, B.F. Santos, “Analyzing Emerging Challenges For Data-Driven Predictive Aircraft Maintenance Using Agent-Based Modeling And Hazard Identification”, Aerospace, 2023, 10 (2), 186.
22. I. Kabashkin, A. Yatskiv, M. Prentkovskis, E. Kabashkina, “Ecosystem Of Aviation Maintenance: Transition From Aircraft Maintenance To Aviation Maintenance 4.0”, Applied Sciences, 2024, 14 (11), 4394.
23. A.D. Kwakye, I.K. Jennions, C.M. Ezhilarasu, “Platform Health Management For Aircraft Maintenance - A Review”, Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering, 2024, 238 (3).
24. S. Fu, N.P. Avdelidis, “Prognostic And Health Management Of Critical Aircraft Systems And Components: An Overview”, Sensors, 2023, 23 (19), 8124.
25. J. Oehling, Using Machine Learning Methods In Airline Flight Data Monitoring To Generate Novel Operational Safety Knowledge, Ph.D. dissertation, Cranfield University, 2019.
26. S.K. Jasra, A. Nichanian, B.F. Santos, “A Comparative Study Of Unsupervised Machine Learning Methods For Anomaly Detection In Flight Data Monitoring”, Aerospace, 2025, 12 (2), 151.
27. N. Cartocci, D.V. Dimarogonas, M.K. Camlibel, “A Comprehensive Case Study Of Data-Driven Methods For Flight Control System Sensor Fault Diagnosis”, Sensors, 2021, 21 (5), 1645.
28. R. Isermann, “Model-Based Fault-Detection And Diagnosis - Status And Applications”, Annual Reviews in Control, 2005, 29 (1), 71-85.
29. V. Venkatasubramanian, R. Rengaswamy, K. Yin, S.N. Kavuri, “A Review Of Process Fault Detection And Diagnosis: Part One: Quantitative Model-Based Methods”, Computers And Chemical Engineering, 2003, 27 (3), 293-311.
30. V. Venkatasubramanian, R. Rengaswamy, S.N. Kavuri, “A Review Of Process Fault Detection And Diagnosis: Part Two: Qualitative Models And Search Strategies”, Computers And Chemical Engineering, 2003, 27 (3), 313-326.
31. V. Venkatasubramanian, R. Rengaswamy, S.N. Kavuri, K. Yin, “A Review Of Process Fault Detection And Diagnosis: Part Three: Process History Based Methods”, Computers And Chemical Engineering, 2003, 27 (3), 327-346.
32. Z. Chen, J. Xu, C. Alippi, S.X. Ding, Y. Shardt, T. Peng, C. Yang, “Graph Neural Network-Based Fault Diagnosis: A Review”, arXiv:2111.08185, 2021.
33. S. Hochreiter, J. Schmidhuber, “Long Short-Term Memory”, Neural Computation, 1997, 9 (8), 1735-1780.
34. K. Cho et al., “Learning Phrase Representations Using RNN Encoder-Decoder For Statistical Machine Translation”, arXiv:1406.1078, 2014.
35. S. Bai, J.Z. Kolter, V. Koltun, “An Empirical Evaluation Of Generic Convolutional And Recurrent Networks For Sequence Modeling”, arXiv:1803.01271, 2018.
36. A. Vaswani et al., “Attention Is All You Need”, Advances in Neural Information Processing Systems, 2017, 30.
37. L. Breiman, “Random Forests”, Machine Learning, 2001, 45, 5-32.
38. J.H. Friedman, “Greedy Function Approximation: A Gradient Boosting Machine”, The Annals of Statistics, 2001, 29 (5), 1189-1232.

39. M. Sakurada, T. Yairi, “Anomaly Detection Using Autoencoders With Nonlinear Dimensionality Reduction”, *Proceedings of the Workshop on Machine Learning for Sensory Data Analysis*, 2014, 4:1-4:11.
40. P.W. Battaglia et al., “Relational Inductive Biases, Deep Learning, And Graph Networks”, *arXiv:1806.01261*, 2018.
41. K.P. Murphy, *Dynamic Bayesian Networks: Representation, Inference And Learning*, Ph.D. dissertation, University of California, Berkeley, 2002.
42. H.B. McMahan, E. Moore, D. Ramage, S. Hampson, B. Aguerre y Arcas, “Communication-Efficient Learning Of Deep Networks From Decentralized Data”, *Proceedings of the International Conference on Artificial Intelligence and Statistics*, 2017, 54, 1273-1282.
43. A. Bemani, N. Ardeshiri, A.H. Balochian, “Aggregation Strategy On Federated Machine Learning Algorithm For Predictive Maintenance”, *Sensors*, 2022, 22 (16), 6252.
44. W. Zhang et al., “Federated Transfer Learning For Remaining Useful Life Predictions Under Heterogeneous Operating Conditions”, *Measurement Science And Technology*, 2025.
45. M.T. Ribeiro, S. Singh, C. Guestrin, “Why Should I Trust You? Explaining The Predictions Of Any Classifier”, *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016, 1135-1144.
46. S.M. Lundberg, S.I. Lee, “A Unified Approach To Interpreting Model Predictions”, *Advances in Neural Information Processing Systems*, 2017, 30.
47. D. Gunning, D. Aha, “DARPA’s Explainable Artificial Intelligence Program”, *AI Magazine*, 2019, 40 (2), 44-58.
48. R. Parasuraman, T.B. Sheridan, C.D. Wickens, “A Model For Types And Levels Of Human Interaction With Automation”, *IEEE Transactions On Systems, Man, And Cybernetics Part A*, 2000, 30 (3), 286-297.
49. J.D. Lee, K.A. See, “Trust In Automation: Designing For Appropriate Reliance”, *Human Factors*, 2004, 46 (1), 50-80.
50. M.R. Endsley, “From Here To Autonomy: Lessons Learned From Human-Automation Research”, *Human Factors*, 2017, 59 (1), 5-27.
51. B. Luettig, C. Torens, V.D. Blondel, “ML Meets Aerospace: Challenges Of Certifying Airborne AI”, *Frontiers in Space Technologies*, 2024, 5, 1475139.
52. A. Degas et al., “A Survey On Artificial Intelligence And Explainable Artificial Intelligence In Air Traffic Management”, *Applied Sciences*, 2022, 12 (3), 1295.
53. A. Nanyonga, B. Turhan, A.N. Ndugga, “A Systematic Review Of Machine Learning Analytic Safety In Aviation”, *Sciences*, 2025, 7 (3), 124.
54. P. Lewis et al., “Retrieval-Augmented Generation For Knowledge-Intensive NLP Tasks”, *Advances in Neural Information Processing Systems*, 2020, 33.
55. R.S. Sutton, A.G. Barto, *Reinforcement Learning: An Introduction*, 2nd ed., MIT Press, Cambridge, MA, USA, 2018.
56. C. Silva, M. Pinheiro, M. Oliveira, “Adaptive Reinforcement Learning For Task Scheduling In Aircraft Predictive Maintenance”, *Scientific Reports*, 2023, 13.
57. A. Pereira, C. Thomas, F. Salgado, “Challenges Of Machine Learning Applied To Safety-Critical Cyber-Physical Systems”, *Machine Learning And Knowledge Extraction*, 2020, 2 (4), 579-602.
58. A.K.S. Jardine, D. Lin, D. Banjevic, “A Review On Machinery Diagnostics And Prognostics Implementing Condition-Based Maintenance”, *Mechanical Systems And Signal Processing*, 2006, 20 (7), 1483-1510.
59. Y. Lei, N. Li, L. Guo, N. Li, T. Yan, J. Lin, “Machinery Health Prognostics: A Systematic Review From Data Acquisition To RUL Prediction”, *Mechanical Systems And Signal Processing*, 2018, 104, 799-834.

APPENDIX




Fig. 4. Continuous Post-Flight Learning and Governance Loop

A. Sample AI Alert Format

```
{
  "alert_trace_id": "ASFN-2026-000001",
  "aircraft_identifier": "AC-EXAMPLE-017",
  "flight_phase": "climb",
  "system_affected": "electrical",
  "risk_severity": "high",
  "confidence_score": 0.87,
  "supporting_sensor_evidence": ["bus_voltage_variance_up", "temperature_residual_persistent"],
  "recommended_inspection": "electrical distribution check on arrival",
  "required_specialist_team": "electrical systems",
  "estimated_urgency": "before next departure",
  "suggested_ground_action": "prepare specialist, diagnostic tooling, and spares",
  "human_validation_status": "pending"
}
```

TABLE XVII. RISK SCORING TABLE

Composite Risk Score	Meaning	Recommended Behavior	Automation
0.00-0.24	Low	Log and trend	
0.25-0.49	Moderate	Advisory only	
0.50-0.74	High	Advisory plus ground pre-notification	
0.75-1.00	Critical	Immediate advisory plus prioritized coordination and approval gate	

TABLE XVIII. EXPERIMENTAL CONFIGURATION

Component	Proposed Configuration
Edge Anomaly Model	GRU or TCN with quantized deployment
Twin Simulation Horizon	5, 15, and 60 minute horizons
Federated Round Frequency	Post-flight or end-of-day aggregation
Alert Calibration	Phase-dependent thresholds with uncertainty adjustment
Explanation Method	SHAP or LIME for tabular models; attention summaries for sequence models
Drift Monitor	Population stability and residual distribution shift

TABLE XIX. REPRODUCIBILITY CHECKLIST

Item	Required Status
Dataset provenance documented	Yes
Train, validation, and test splits disclosed	Yes

“AI-SAFE FlightNet: A Next-Generation Artificial Intelligence and Automation Framework for Predictive Flight Safety, Aircraft Health Intelligence, And Ground Operations Coordination”

Item	Required Status
Synthetic telemetry generator described	Yes
Hyperparameters documented	Yes
Random seeds reported	Yes
Evaluation metrics formally defined	Yes
Ablation conditions specified	Yes
Code release plan stated	Yes
Safety and cyber assumptions stated	Yes

TABLE XX. AVIATION AI GOVERNANCE CHECKLIST

Governance Control	Required
Human authority boundary defined	Yes
Approved evidence sources listed	Yes
Model versioning active	Yes
Data lineage auditable	Yes
Drift monitoring deployed	Yes
Cybersecurity controls implemented	Yes
Override logging enabled	Yes
Post-flight learning gated by validation	Yes
Safety case artifacts maintained	Yes
Role-specific explanation policy defined	Yes