Journal of Industrial Policy and Technology Management, 9(1), 2026, 53-80

Federated Fleet Learning with Explainable AI: A Privacy-Preserving Architecture for Aviation Predictive Maintenance

Akif Emrah Büyüksomer[a], **Tuncel Öz**[b]

a) İstanbul Ticaret University, b) İstanbul Ticaret University, toz@ticaret.edu.tr,

10.5281/zenodo.20517392

Received: Feb 25, 2026 Accepted: Mar 23, 2026 Published: Jun 05, 2026

Abstract: Engine maintenance stands at the centre of operational sustainability and cost governance in civil aviation, yet the sector confronts a paradox in data governance. Unscheduled failures impose severe financial penalties and safety hazards, but carriers remain reluctant to share operational data—treating it as proprietary intelligence—and competitive pressures suppress horizontal collaboration. Centralised machine-learning paradigms that require raw data transfer are structurally incompatible with these constraints. This paper presents a comprehensive literature review designed to address this deadlock from the standpoint of information systems and technology management. It conceptualises a Federated Fleet Learning architecture that enables a shared corporate intelligence network while preserving Data Sovereignty. Three interlocking components are examined as pillars of a Decision Support System: the 1DCNN-BiLSTM hybrid, suited to the constrained processing capacity of Aircraft Edge Units; the FedProx optimisation algorithm for managing statistical heterogeneity across disparate fleets; and asynchronous LEO satellite strategies for mitigating network latency. SHAP-based Explainable AI (XAI) integration is scrutinised within a technology governance frame to meet FAA and EASA DO-178C certification demands and to ensure algorithmic accountability. A noteworthy finding emerges: fleets with disjoint failure profiles—one versed solely in turbine degradation, another in compressor faults—can illuminate each other's corporate “Blind Spots” through Cross-Fleet Knowledge Transfer without any raw data exchange. The study amalgamates algorithmic feasibility with industrial policy, laying a theoretical foundation for a privacy-preserving digital transformation paradigm in aviation.

Keyword: AI Technologies and Management, Data Science and Business Analytics, Digital Transformation, Decision Support Systems, Federated Fleet Learning, Explainable AI (XAI)

JEL Classification: O32, L93, M15, C45

1. Introduction and Conceptual Framework

Civil aviation constitutes one of the most technologically complex and safety-critical sectors in the contemporary global economy. Among the foremost operational challenges that carriers face daily is the effective governance of engine maintenance, repair, and overhaul (MRO) processes. The magnitude of this challenge is considerable: according to the International Air Transport Association, global airlines devoted approximately USD 93.9 billion to MRO activities in 2023, a sum equivalent to eleven per cent of total operating costs (IATA, 2024). This vast ecosystem attests to how deeply maintenance processes are interwoven with airline operations and logistics performance (Demirbilek, Öz & Fidan, 2018). Oliver Wyman's (2025) projections paint an even starker picture, reporting that global MRO expenditures have reached USD 114 billion while the active fleet has expanded to roughly 29,000 aircraft.

When an engine is incapacitated by an unforeseen failure—whether in flight or on the tarmac—the event precipitates what the literature designates an "Aircraft on Ground" (AOG) incident. The annual cost of such groundings alone exceeds USD 50 billion. A question follows naturally: is the ability to anticipate failures with high confidence purely an engineering concern, or does it represent a strategic economic imperative for airlines?

Conventional maintenance philosophies rest on either reactive interventions indexed to fixed time intervals or flight-hour thresholds, or on static scheduled overhauls (Mobley, 2002; Dragomir et al., 2009). Although these strategies served the industry for decades, they no longer satisfy the dynamic, low-fault-tolerance demands of modern aviation. The Predictive Maintenance paradigm emerged to overcome this limitation. By estimating a component's Remaining Useful Life (RUL) from its real-time condition data, the approach detects incipient problems far earlier and curtails unplanned removals (Fu & Avdelidis, 2023; Achouch et al., 2023). Given that a modern commercial aircraft can generate hundreds of terabytes of sensor data even on a short-haul flight (Badea et al., 2018), machine-learning algorithms form the backbone of this paradigm. The state of the art and future trajectories of AI-driven Predictive Maintenance in aviation have been surveyed extensively by Uçar, Karaköse and Kırımça (2024) and by Stanton et al. (2023).

The sector, however, is ensnared in a paradox born of its own competitive dynamics. Unplanned failures generate trillion-scale liabilities, yet carriers remain unwilling to share data. Operational records are treated as trade secrets; rivalry within the Airbus and Boeing ecosystems erects barriers to lateral cooperation. Centralised machine-learning approaches that mandate raw-data pooling are fundamentally at odds with these privacy expectations. This gridlock produces what the present study terms the "Blind Spot" problem: a fleet that collects data in isolation remains effectively blind to failure patterns encountered by foreign fleets. The strategic management of digital transformation in enterprises necessitates a redesign of data-driven decision-making and cooperation models, a pressure that applies with equal force to airline fleets (Kumru & Kasımoğlu, 2022).

Federated Learning (FL) enters the stage at this juncture (McMahan et al., 2017). By confining data to local devices and transmitting only trained model weights to a central server, this architecture offers an innovative pathway for constructing collective intelligence without surrendering Data Sovereignty (Kairouz et al., 2021; Abdulrahman et al., 2021). Recent studies on RUL estimation and fault diagnosis in privacy-preserving settings (Zhang & Li, 2022; Bemani & Björsell, 2023) have established the industrial viability of this paradigm.

The present literature review seeks to synthesise the Predictive Maintenance concept for civil-aviation engines along eleven thematic axes: the algorithmic foundations of federated learning, deep-learning architectures for time-series forecasting, turbofan engine prognostics, anomaly detection, gradient-boosting methods, Explainable AI (XAI), aviation regulations, LEO satellite communications, data-privacy and security mechanisms, edge-computing applications, and hybrid modelling approaches. The evolution of Decision Support Systems from model-driven to knowledge-driven architectures signals a critical paradigm shift in the integration of AI-based analytics into corporate decision-making processes (Hızıroğlu, Pişirgen, Özcan & İlter, 2022).

2. Research Methodology

This literature review was conducted to map, in a systematic fashion, the extant body of knowledge at the intersection of Predictive Maintenance for aviation fleets, federated learning, and Explainable AI. Six databases served as the primary source pools: Web of Science (WoS), Scopus, IEEE Xplore, ScienceDirect (Elsevier), SpringerLink, and arXiv. The search strategy employed the following keyword combinations: “federated learning” AND (“predictive maintenance” OR “remaining useful life” OR “prognostics”); “explainable AI” AND (“aviation” OR “turbofan” OR “aircraft engine”); “LEO satellite” AND (“federated learning” OR “aviation connectivity”); “N-CMAPSS” OR “C-MAPSS” AND (“deep learning” OR “RUL prediction”). For the Turkish-language literature, the terms “federe öğrenme” and “öngörücü bakım” were additionally searched.

Inclusion criteria were defined as follows: (i) peer-reviewed journal articles, conference proceedings, and selected preprints published between 2017 and 2025; (ii) studies written in English or Turkish; (iii) research addressing at least one of the following topics directly: Predictive Maintenance, federated learning, Explainable AI, turbofan prognostics, or aviation communication infrastructure. Exclusion criteria comprised: (i) studies whose full texts were inaccessible; (ii) letters to editors, brief notes, or poster abstracts; (iii) general-purpose machine-learning studies that could not be directly linked to aviation or industrial maintenance contexts; (iv) where multiple versions of the same study existed, only the most recent was retained.

The initial screening identified over 450 candidate sources. After a title-and-abstract review, this figure was narrowed to 280 studies; a rigorous full-text evaluation subsequently distilled the final bibliography to upward of 170 references. Snowballing was applied throughout: both backward tracking through reference lists and forward tracking through citing articles expanded the review's scope. The resulting corpus was classified under eleven thematic sections, and in each section the strengths, limitations, and research gaps of the existing literature were appraised from a critical standpoint.

3. Federated Learning Fundamentals and Comparative Analysis of Algorithmic Advances

Federated learning is a paradigmatic approach that enables the training of a single model from decentralised data sources while preserving data privacy in absolute terms. The founding algorithm, FedAvg (Federated Averaging), rests on the principle that each client applies stochastic gradient descent on its own local data and the resulting updates are iteratively averaged at a central server (McMahan et al., 2017). Its simplicity and low computational overhead render FedAvg attractive for industrial deployments. The fundamentals and challenges of federated learning have been surveyed comprehensively by Beltrán et al. (2023) and Li et al. (2020a).

FedAvg's architecture, however, exhibits serious vulnerabilities when confronted with the Non-Independent and Identically Distributed (Non-IID) heterogeneous datasets that are endemic to industrial environments such as aviation. Data imbalances across airline fleets cause local models to diverge progressively from the global model, culminating in "client drift." FedProx was developed to counteract this phenomenon: it appends a proximal penalty term to the local update objective, thereby mathematically constraining the extent to which client models may deviate from the global model (Li et al., 2020b). The principal advantage of FedProx is its capacity to maintain stable convergence even under pronounced data heterogeneity; the practical difficulty lies in the need to tune the penalty coefficient (μ) with precision.

SCAFFOLD offers a different analytical lens on client drift. It employs control variates to rectify the discrepancy between local and global update directions (Karimireddy et al., 2020). By storing control variances at each client, SCAFFOLD corrects the update trajectory and achieves heterogeneity-agnostic convergence. The trade-off is a substantial increase in memory and processor requirements per client—a disadvantage in aviation settings where edge devices are resource-constrained. Beyond these three core algorithms, alternatives such as FedNova, which normalises updates to balance disparate computational capacities (Wang et al., 2020), and FedAdam/FedYogi, which leverage adaptive learning rates (Reddi et al., 2021), populate the literature.

The aviation-specific strengths and weaknesses of these three federated optimisation strategies are synthesised in Table 1.

Table 1. *Comparative Analysis of Federated Learning Algorithms (Non-IID Data Scenarios)*

Algorithm	**Core Approach**	**Advantages**	**Disadvantages**	**Aviation Prognostics**
FedAvg	Averaging of local gradients	Low computational complexity; straightforward integration (McMahan et al., 2017)	Severe client drift under Non-IID fleets; unstable convergence	Up to 94% accuracy in homogeneous (IID) fleets; rapid degradation in real-world heterogeneous settings
FedProx	Proximal penalty term added to local loss function	High stability under heterogeneous fault data; prevents client drift (Li et al., 2020b)	Initial-round oscillation tendency; sensitive µ tuning required	Most stable architecture for cross-fleet transfer across disparate motor fault types (Non-IID)
SCAFFOLD	Update-direction correction via control variates	Global convergence in fewer communication rounds (Karimireddy et al., 2020)	Substantially increased memory and edge-processor demands per client	Imposes high system load in aviation contexts where edge devices are constrained

As Table 1 makes plain, no single algorithm dominates across all scenarios. Algorithm selection in aviation federated learning must therefore be calibrated to the fleet's Non-IID severity, the processing capacity of Aircraft Edge Units, and the available satellite bandwidth. Communication costs constitute a critical bottleneck in satellite networks with narrow bandwidth (Konečný et al., 2016). When the irregular connectivity windows of airborne platforms are factored in, conventional synchronous federated approaches lose viability altogether. In synchronous designs, the server must await updates from every aircraft, meaning that the slowest connection blocks the entire fleet's training cycle—the so-called straggler effect (Öz, 2020).

Asynchronous Federated Learning (AFL) was developed to circumvent this constraint. Recent work demonstrates that FedFa, a fully asynchronous training paradigm (Xu et al., 2024), delivers strong performance in overcoming these communication bottlenecks.

4. Deep Learning Architectures for Time-Series Prediction: A Comparative Examination

Sensor readings harvested from turbofan systems—temperature fluctuations, turbine rotational speeds, fuel-flow rates—constitute multivariate time-series data that characterise engine health. Detecting sequential patterns within these complex datasets and predicting Remaining Useful Life (RUL) with high accuracy falls squarely within the remit of deep neural networks (Zhao et al., 2019; Serradilla et al., 2022).

Long Short-Term Memory (LSTM) networks, widely regarded as a watershed in time-series modelling, resolved the vanishing-gradient problem of classical recurrent networks through their forget-gate and input-gate cell architecture (Hochreiter & Schmidhuber, 1997). Classical LSTMs, however, read time in only one direction—past to future. This unidirectional structure can prove insufficient for capturing the contextual relationships inherent in gradual degradation. Bidirectional LSTM (BiLSTM) overcomes this limitation by reading the time series simultaneously in both forward and reverse directions, modelling both the antecedent conditions and the prospective degradation trajectory of a fault in a single pass (Sherifi, 2024). The result is a marked accuracy gain over unidirectional architectures in detecting progressive damage patterns.

One of the most consequential advances in deep learning has been the advent of hybrid architectures that fuse spatial feature extraction with temporal sequencing. Structures coupling 1D Convolutional Neural Networks (1D-CNN) with BiLSTM extract local inter-sensor correlations with high fidelity, while the BiLSTM component processes these condensed features along the temporal axis to shape the RUL estimate (Wahid et al., 2022; Mo et al., 2023). Recent years have confirmed that such hybrid architectures deliver highly competitive results on NASA C-MAPSS benchmarks (Yu et al., 2025). In parallel, Transformer-based models—which abandon recurrent structures in favour of parallel processing—have entered the race. Their self-attention mechanism captures long-range degradation patterns directly (Vaswani et al., 2017), and their application to RUL prediction is expanding rapidly (Wen et al., 2022).

The prediction performance and hardware demands of these architectures on NASA datasets are compared in Table 2.

Table 2. RUL Prediction Performance of Time-Series Deep Learning Architectures

Model Architecture	Sensor Feature Extraction Strategy	Prediction Accuracy Assessment	Hardware & Processing Cost (On-Board Use)
Conventional LSTM	Unidirectional sequential processing; limited spatial awareness	Relatively high error tendency on N-CMAPSS (El-Rashidy et al., 2022)	Low computational cost. Readily trainable on constrained hardware.
1DCNN-BiLSTM	CNN-based spatial filtering + BiLSTM bidirectional temporal context	Low error rates. Strong capability for modelling compound faults (Wahid et al., 2022; Yu et al., 2025)	Moderate. Offers an optimal architectural balance for sensor fusion and federated distribution.
Transformer	Parallel self-attention matrices; no recurrence	Superior RUL performance on long-range degradation patterns (Wen et al., 2022; Guo et al., 2024)	Very high. Massive hyper-parameter space strains Aircraft Edge Unit processors.
FTT-GRU (Hybrid)	Fast Fourier Transform-simplified attention	Error rates comparable to conventional models at very low latency (Deng & Zhou, 2024)	Low latency; compatible with real-time flight telemetry.

Where maximum accuracy is the sole objective, Transformer architectures take the lead. Yet a caveat cannot be ignored: once the constrained processing capacity of Aircraft Edge Units is factored in, hybrid structures such as 1DCNN-BiLSTM offer an optimal engineering trade-off between accuracy and computational cost for real-world federated deployments. Considering that even a one-percentage-point error reduction in long-range degradation patterns can be consequential for flight safety, the choice of architecture is, at its core, a strategic decision.

5. Turbofan Engine Prognostics and Overcoming Sectoral "Blind Spots"

Turbine engines—systems in which tens of thousands of moving parts operate under extreme thermal, mechanical, and pressure stresses—are the most valuable propulsive assets in civil aviation. Prognostics, the discipline of analysing latent fault signatures and predicting run-to-failure time, is indispensable for aviation safety (Lei et al., 2021; Fu & Avdelidis, 2023).

For years the original C-MAPSS dataset produced by NASA served as the primary benchmark. Its fixed-altitude and steady-regime scenarios, however, are limited in representing the climb, cruise, and descent cycles of modern aircraft (Zhang et al., 2020; Ajay et al., 2023). The New C-MAPSS (N-CMAPSS) dataset, released in 2021 (Arias Chao et al., 2021), broke new ground for prognostic research by offering simulated realistic flight profiles and multiple failure scenarios that far surpass the original's richness—dynamic altitude changes, multi-regime operations, and highly intricate degradation mechanisms.

Of particular significance is its capacity to test "Blind Spot" scenarios. Building on the literature reviewed here, a strategic matching methodology using N-CMAPSS subsets is proposed to address data asymmetry and to guide future simulations: the integration of sub-datasets DS03 and DS04. DS03 simulates a seven-aircraft (fourteen-engine) fleet conversant with High- and Low-Pressure Turbine (HPT and LPT) failures in the engine's hot section, whereas DS04 represents a five-aircraft (ten-engine) fleet whose failures are confined exclusively to the cold-section Fan and Low-Pressure Compressor (LPC). The degradation dynamics of hot-section turbines and the fault signatures of cold-section compressors are entirely distinct.

Within a decentralised federated network, the ability of the DS03 and DS04 fleets to learn each other's failure patterns without exchanging raw data carries the potential to constitute one of the clearest demonstrations of horizontal cooperation in aviation history (Landau et al., 2026). In this federated scenario, each fleet's task of modelling failure types it has never encountered reduces, ultimately, to a domain-adaptation problem; transfer-learning-based knowledge exchange between subsets with different fault dynamics substantially enhances the model's generalisation capacity.

One of the most consequential theoretical contributions that this review offers to the literature is the methodological recommendation that Engine 15 in DS03 be isolated entirely from the training process and reserved as a "Zero-Shot / Blind Spot" validation set. This protocol would permit future empirical studies to demonstrate, with full transparency, the extent to which a model can diagnose a hitherto unseen failure type by leveraging experience acquired through the federated network. Such a validation design also eliminates the risk of data leakage, thereby reinforcing the scientific credibility of the findings (Öz et al. 2019).

6. Anomaly Detection and Fault Diagnosis Processes

Anomaly detection and fault diagnosis, two closely related yet operationally distinct stages, underpin Predictive Maintenance. Anomaly detection aims to capture deviations from normal operating patterns in sensor data; fault diagnosis seeks to determine the origin and type of that deviation.

One-Class Support Vector Machines (One-Class SVM) were historically favoured for these tasks, yet they present serious scalability difficulties in high-dimensional sensor data (Fernandes et al., 2022). Deep-learning-based Autoencoders have supplanted them. Park et al. (2019) proposed a hybrid anomaly-detection framework combining an autoencoder with LSTM. Ahmed et al. (2023) developed an intelligent anomaly-detection system based on feature autoencoders for industrial machinery. In these approaches the autoencoder learns the engine's "healthy" operating profile and triggers an anomaly signal when the reconstruction error exceeds a defined threshold. Qiu et al. (2023) furnished a comprehensive survey of deep learning techniques in intelligent fault diagnosis and prognosis. Son et al. (2023) achieved simultaneous detection of multiple fault types through a convolutional-neural-network-based multi-output classification model.

7. XGBoost and Gradient-Boosting Tree Architectures

Gradient boosting is an effective ensemble-learning approach in which weak learners are chained so that each corrects the errors of its predecessor. XGBoost, its most powerful representative, has gained widespread acceptance in prognostics owing to its built-in regularisation mechanisms and its resilience to missing data. The role of feature engineering and aggregated feature importance in aircraft-engine RUL prediction has been examined by Alomari et al. (2023) in Scientific Reports.

What distinguishes XGBoost in aviation engineering is its native interpretability. Unlike certain deep networks that cannot explain their reasoning, XGBoost quantitatively scores the contribution of each sensor to a failure prediction (Feature Importance). This capability is amplified when XGBoost is hybridised with Bayesian networks; Velasco-Loera et al. (2025) proposed an interpretable hybrid fault-prediction framework in Applied Sciences. In some contemporary pipelines XGBoost is deployed not as the primary prediction model but as a highly adept feature selector that eliminates noise-generating sensors, thereby shielding complex deep-learning models from overfitting (Liu et al., 2022).

LightGBM, with its histogram-based learning structure, is another powerful gradient-boosting algorithm employed in this domain (Li et al., 2018; Li et al., 2024b). Deep CNN and LightGBM combinations for RUL prediction have been investigated by Ma et al. (2021). Nemat Saberi et al. (2022) developed a LightGBM-based method for rotating-machinery fault diagnosis under non-stationary conditions.

Training these algorithms on a federated network where data is never centralised poses substantial methodological challenges. XGBoost inherently requires the global distribution of data for its branching decisions. To circumvent this bottleneck, "Secret Sharing"-based Federated XGBoost algorithms have been developed—approaches that share only encrypted distributional statistics and never pool raw data (Xie et al., 2022). Scalable and secure Federated XGBoost was presented by Nguyen et al. (2023) at ICASSP 2023. Federated fuzzy regression trees were examined by Corcuera Bárcena et al. (2025) in Information Fusion.

8. Explainable AI (XAI): Making the Black Box Transparent in Aviation

In a sector where safety requirements are paramount, it is insufficient for an AI system merely to recommend that an engine be taken in for maintenance after fifty flight cycles. If the system cannot transparently articulate the logical grounds for its recommendation, operational acceptance is precluded. Explainable AI (XAI) is no longer a technological luxury; it is a regulatory obligation anticipated by international aviation authorities (Cummins et al., 2024).

The literature offers several approaches designed to illuminate models' decision logic. Fan et al. (2021) provided a comprehensive review of interpretability techniques for artificial neural networks. The current state and future directions of XAI for time-series classification were discussed by Theissler et al. (2022) in IEEE Access. In AI-based medical Decision Support Systems developed across various regions, the necessity for clinicians to interpret model predictions with confidence has been stressed in numerous studies (Karagül Yıldız, Yurtay & Öneç, 2021). The two most prominent methods are SHAP (SHapley Additive exPlanations), grounded in game theory, and LIME (Local Interpretable Model-agnostic Explanations), which operates locally (Lundberg & Lee, 2017; Salih et al., 2023). The functions and relative merits of these two methods in the context of aviation prognostics are compared in Table 3.

Table 3. *Comparison of Explainable AI (XAI) Methods in Aviation Engine Maintenance*

Criterion / Feature	**SHAP (SHapley Additive exPlanations)**	**LIME (Local Interpretable Model-agnostic Explanations)**
Core Mechanism	Game theory (distributing marginal impact via Shapley values) (Lundberg & Lee, 2017)	Building a local linear model by generating synthetic data around a prediction (Salih et al., 2023)
Explanation Scope	Both local and global analyses (Hong et al., 2020)	Local analyses only
Mathematical Consistency	Consistent. Always identifies the same feature for identical sensor input.	May exhibit instability due to dependence on stochastic synthetic generation.
Edge-Computing Load	Computationally heavier and more time-consuming.	Produces a single explanation rapidly.
Aviation Regulatory Compliance	Deemed more reliable for safety-critical processes owing to its consistency.	Risk of producing contradictory results may pose a handicap in certification.

Although LIME delivers rapid explanations, its mathematical reliance on randomness can produce inconsistent outputs. SHAP's consistent structure—quantifying each sensor's contribution to the prediction with precision—has earned it recognition as the most reliable tool for overcoming the black-box problem among civil-aviation regulators (Pan et al., 2023). The performance of XAI methods in asset-failure prediction has been evaluated experimentally by Jakubowski et al. (2022).

The concept of “Federated Explainable AI” (fXAI) was introduced by Kusiak (2024) in the International Journal of Production Research and discussed from a digital-manufacturing perspective. Practical applications of this concept were explored by Alshkeili et al. (2025) in a privacy-preserving interpretable federated learning model. Lopez-Ramos et al. (2024) examined the interplay between federated learning and XAI through a scoping review. Amato and Branco (2025) proposed SemFedXAI, a semantic framework for explainable federated learning in healthcare. Enabling trustworthy federated learning in Industrial IoT—bridging interpretability and robustness—was addressed by Jagatheesaperumal et al. (2024).

9. Aviation Regulations, Predictive Maintenance, and DO-178C Compliance Challenges

The deployment of new technologies in aviation is subject to exceedingly stringent certification processes. The Federal Aviation Administration (FAA) and the European Union Aviation Safety Agency (EASA) are the primary standard-setting bodies. DO-178C, the sector’s baseline standard, mandates that software behaviour be one-hundred-per-cent deterministic (FAA, 2011). Self-learning machine-learning models that evolve on the basis of data are, by construction, at odds with this deterministic certification philosophy.

To reconcile this tension, EASA published its “Artificial Intelligence Roadmap 2.0,” introducing a “human-centric” AI concept for aviation (EASA, 2023). Under EASA’s framework, AI should be positioned not as an autonomous actor intervening directly in systems but as an assistant supporting pilots and maintenance engineers. Automation levels are tiered as Level 1 (Human Support) and Level 2 (Human-Machine Collaboration), with transparency, traceability, data privacy, and explainability established as core principles (EASA, 2020, 2023). EASA’s W-shaped development model extends the traditional V-model with a learning-assurance step, furnishing a unique framework for AI certification in aviation.

The aviation industry’s technological complexity, safety requirements, and economic scale create an exceptionally fertile yet challenging terrain for Predictive Maintenance. Digital twin technology emerges as a significant complementary element: Xiong et al. (2021) examined digital-twin-driven intelligent Predictive Maintenance for aero-engines, while Bisanti et al. (2024) reviewed the role of digital twins in aircraft maintenance and operations through a systematic literature review. The challenges and opportunities of condition-based maintenance (CBM) in aviation were discussed by Verhagen et al. (2023); Eliaz et al. (2022) offered a systematic review for defence fixed-wing aircraft. AI-driven fault detection and maintenance optimisation from an aviation technical-support perspective were assessed by Kabashkin et al. (2025). The interpretability of machine-learning models in aircraft-engine Predictive Maintenance was addressed in a systematic review by Al Hasib et al. (2023).

In DAL A or DAL B categories—where software errors can directly jeopardise flight safety—the approach to AI-assisted code is highly cautious. When an AI algorithm issues a maintenance recommendation, the reasoning behind that decision must be traceable through SHAP-like methods and verifiable by human engineers. No algorithm lacking accountability and whose decision processes cannot be articulated appears likely to pass civil-aviation certification (Uçar et al., 2024; Scarselli & Nicassio, 2025).

10. LEO Satellite Communication: An Asynchronous Network Infrastructure from Space

For modern aircraft to participate in a synchronous federated-learning network with ground stations thousands of kilometres away, an uninterrupted, low-latency communication infrastructure is essential. Legacy Geostationary-Orbit (GEO) satellites, positioned roughly 36,000 km from Earth, introduce operationally unacceptable delays of 600–800 ms, severely impeding model synchronisation (Abdallah & Psomas, 2021). If each communication round in federated training completes with a latency of several hundred milliseconds, global-model convergence can stretch from hours to days.

The aviation satellite-communication market is undergoing a profound technological transformation. SpaceX's Starlink constellation, comprising over 9,400 active satellites and inter-satellite laser crosslinks, enables intercontinental data transfer without descending to the surface (SpaceX, 2025). Michel et al. (2022) conducted a comprehensive performance analysis of Starlink at IMC 2022. Mohan et al. (2024), drawing on 19.2 million speed tests, demonstrated that Starlink achieves 20–40 ms latency and download speeds exceeding 100 Mbps across 34 countries. These figures point to an infrastructure ideally suited for asynchronous federated weight updates.

OneWeb/Eutelsat (now a combined entity) operates a constellation of approximately 648 satellites at an altitude of 1,200 km, offering 40–70 ms latency (Eutelsat Communications, 2023). SES has carved out a position in the commercial-aviation market with multi-orbit (LEO+GEO) hybrid solutions (SES S.A, 2024). The hardware-level implications of different satellite architectures for federated-learning training are compared in Table 4.

Table 4. Satellite Architecture Comparison for Aviation Connectivity and Federated Learning

Feature / Metric	Starlink (Pure LEO)	OneWeb/Eutelsat (LEO+GEO)	GEO Legacy Satellites
Orbital Altitude	550 km, high density (SpaceX, 2025)	1,200 km, multi-orbit (Eutelsat, 2023)	36,000 km
Connection Latency	20–40 ms (Michel et al., 2022; Starlink, 2024)	40–70 ms (Eutelsat, 2023)	600–800 ms
Network Infrastructure	Laser crosslinks enable direct inter-satellite transfer (SpaceX, 2024)	Requires descent to Satellite Network Portal (Eutelsat, 2023)	Dependent on single ground station
Federated Learning Impact	Ideal low-latency infrastructure for asynchronous weight updates	High redundancy; connectivity fluctuations during handovers	Inoperable for real-time FL

The literature advocates a "Dual-Vendor LEO Strategy" to mitigate the service-interruption risk inherent in reliance on a single network provider. Under such an architecture, the on-board system can switch rapidly from a congested Starlink link to the OneWeb/SES network, guaranteeing the secure transmission of federated model gradients to the aggregation server. This redundancy is especially critical on transoceanic routes, where the probability of falling outside a single provider's coverage footprint is elevated.

11. Data Privacy, Cybersecurity, and Regulatory Challenges

Federated learning's architecture confers a design-level privacy advantage by precluding raw-data sharing. Sharing only model weights and gradients does not, however, eliminate security risks in their entirety. Malicious actors can reverse-engineer these numerical gradients through Gradient Inversion Attacks, reconstructing original sensor data with considerable fidelity (Geng et al., 2023; Huang et al., 2021). Recent frameworks such as HyperFL offer robust defence mechanisms against these attacks via hypernetwork-based federated learning (Guo et al., 2025).

The principal countermeasure is the Differential Privacy (DP) shield. DP injects mathematically calibrated noise into the updates transmitted by aircraft, rendering the original data unrecoverable even if gradients are inverted (El Ouadrhiri & Abdelhadi, 2022; Baraheem & Yao, 2022). A delicate balance obtains: excessive noise degrades model accuracy, while insufficient noise compromises privacy. Secure Multi-Party Computation (SMPC) protocols guarantee that gradients are never shared in plaintext (Bonawitz et al., 2017).

A further major threat arises from "Byzantine Clients"—nodes that have been compromised or whose sensors are faulty—which transmit intentionally corrupted gradients to sabotage the global model (Xia et al., 2025). Server-side robust aggregation algorithms such as Trimmed Mean and Krum successfully isolate these outlier updates (Xu et al., 2025; Zuo et al., 2025). A comprehensive taxonomy of poisoning attacks in federated learning was presented by Xia et al. (2023) in IEEE Access.

12. Federated Learning Applications and Edge Computing in Aviation Predictive Maintenance

The theoretical foundations of federated learning were laid in the preceding sections; the true value of a technology, however, materialises in its practical applications. The opportunities and challenges that federated learning presents for the Industrial Internet of Things (IIoT) have been examined systematically by Nguyen et al. (2021) in both the IEEE Internet of Things Journal and IEEE Communications Surveys and Tutorials.

Direct applications of federated learning to aviation Predictive Maintenance are proliferating. Barbosa et al. (2025) evaluated federated machine learning on the C-MAPSS dataset, demonstrating that decentralised training can yield results approaching those of centralised training. Landau et al. (2026) modelled a multi-airline scenario on N-CMAPSS, proposing a federated learning framework for collaborative RUL prognostics. Kamei and Taghipour (2023) compared centralised and decentralised federated learning using the Transformer architecture. Arunan et al. (2023) contributed to industrial health prognostics on heterogeneous edge devices through matched feature extraction. Predictive maintenance and anomaly detection using distribution shifts in time-series data were examined comprehensively by Ahn et al. (2023).

The edge-computing dimension is equally determinative for operationalisation. The constrained processing capacity of Aircraft Edge Units mandates model compression and communication-cost reduction. Federated learning in cloud-edge collaborative architectures was surveyed by Bao and Guo (2022), who proposed strategies for model downsizing and efficient distributed training given edge devices' limited processing power. Liu et al. (2024) proposed an edge-computing-oriented federated learning framework. Low-latency collaborative Predictive Maintenance through Over-the-Air FL in noisy industrial environments was investigated by Karami et al. (2023). Von Wahl et al. (2024) presented an asynchronous federated learning approach aware of data disparity and temporal unavailability at the AAAI Conference, addressing challenges specific to transportation fleets.

Federated learning's aviation applications beyond Predictive Maintenance are also expanding. The reinforcement-learning-based framework FedPreM was proposed by Yang et al. (2023) at IEEE GLOBECOM. Privacy-aware resource sharing in cross-device federated model training was addressed by Bharti and McGibney (2021); federated generative models in industrial environments were evaluated by Milasheuski et al. (2024). Foundation-model and federated-learning integration for IoT-based aircraft health monitoring was proposed by Kabashkin (2024) in Mathematics.

13. Hybrid Models, Domain Adaptation, and Learning from Scarce Data

Given the complex dynamics of Predictive Maintenance, no single machine-learning algorithm can consistently deliver peak performance across all operational scenarios. Ensemble Learning and hybrid approaches—where one model's weakness is compensated by another's strength—hold the key to optimal prediction accuracy. Integrating the LSTM architecture with classification-robust algorithms such as XGBoost or AdaBoost minimises error margins and yields reliable aviation-maintenance outcomes (Tian & Wang, 2020; Huang et al., 2023). Ensemble neural networks for RUL prediction under uncertainty were investigated by Biswas et al. (2023); a hybrid ensemble deep-learning approach for battery RUL prediction was presented by Xu et al. (2023b) in IEEE/CAA Journal of Automatica Sinica. Transfer-learning techniques that enable knowledge transfer across different operating conditions enhance model generalisability in environments with disparate data distributions (Pan & Yang, 2010). Data-driven Predictive Maintenance literature broadly reports that ensemble and hybrid modelling strategies improve RUL prediction accuracy (Zhang, Yang & Wang, 2019).

Physics-Informed Neural Networks (PINNs) address the risk that purely data-driven models may produce physically implausible outputs. By embedding the thermodynamic equations of a turbofan engine as a constraint in the model's loss function, PINNs prevent predictions that violate thermodynamic laws and enable consistent learning with less data (Wen et al., 2023; Kim et al., 2022a). Scalable data-transformation models for PINNs in digital-twin-enabled PHM applications were examined by Barimah et al. (2025) in Computers. Chen et al. (2020) achieved bearing RUL prediction through an attention-mechanism-based deep-learning method.

Because actual engines in aviation are prevented from reaching the run-to-failure stage by flight-safety protocols, a substantial "Small Data Problem" persists (Li et al., 2024a). In environments where labelled failures are exceedingly rare, Transfer Learning and Domain Adaptation techniques are critical (Chen et al., 2023c). Transferring the weights trained on a source fleet to a target fleet compensates for data scarcity and enables domain-agnostic model operation (Li et al., 2024a; Ding et al., 2022). This mechanism directly parallels the DS03/DS04 "Blind Spot" scenario discussed in Section 5: a fleet conversant with hot-section failures learning cold-section faults through the federated network is, at its core, a domain-adaptation problem. A comprehensive review of deep domain adaptation for turbofan RUL prediction was provided by Landau et al. (2026); Transformer-based domain-adaptive RUL prediction was addressed by Li, Zhang, Ding and Sun (2022b) in IEEE Transactions on Instrumentation and Measurement.

The intersection of federated learning and transfer learning is growing richer. Privacy-preserving federated transfer learning for machinery fault diagnosis was examined by Zhang and Li (2022) in Structural Health Monitoring. Label-synchronisation strategies for hybrid federated learning were explored by Llasag Rosero et al. (2025) in Reliability Engineering and System Safety; strategies for overcoming data scarcity, imbalance, and feature-selection challenges in machine-learning models were addressed by Hakami (2024) in Scientific Reports.

14. Comparative Literature Analysis

A systematic comparison of the field's seminal studies is essential for a holistic appraisal of the existing literature. Table 5 compares federated learning applications in aviation Predictive Maintenance on an author-by-author basis.

Table 5. *Comparative Analysis of Federated Learning Studies in Aviation Predictive Maintenance*

Author(s) & Year	**Dataset**	**FL Algorithm**	**DL Architecture**	**XAI**	**Key Contribution**	**Limitations**
Bemani & Björsell (2023)	C-MAPSS	FedAvg	FedSVM, FedLSTM	None	First aviation FL framework	Single dataset; no Non-IID scenario
Zhou et al. (2023)	Custom industrial	FedAvg	CNN	None	Privacy-preserving fault diagnosis	Limited engine-type diversity
Kamei & Taghipour (2023)	C-MAPSS	FedAvg, FedProx	Transformer	None	Centralised vs. decentralised FL comparison	N-CMAPSS not evaluated
Ahn et al. (2023)	C-MAPSS	FedAvg	LSTM	None	Time-series distribution-shift analysis	No asynchronous scenario
Landau et al. (2026)	N-CMAPSS	FedAvg + robust aggregation	CNN-LSTM	None	Collaborative RUL prognostics with decentralised validation	No XAI or LEO integration

Author(s) & Year	**Dataset**	**FL Algorithm**	**DL Architecture**	**XAI**	**Key Contribution**	**Limitations**
Barbosa et al. (2025)	C-MAPSS	FedAvg	MLP, RF	None	Non-linear federated ML achieves near-centralised results	Limited DL architecture
Alshkeili et al. (2025)	General	FL + DP	DNN	SHAP	Privacy-preserving explainable FL model	Not aviation-specific
Conceptualised Framework	**N-CMAPSS (DS03+DS04)**	**FedProx**	**1DCNN-BiLSTM**	**SHAP**	**Dual LEO + Async FL + XAI + Blind Spot Transfer**	**Simulation-based; empirical validation pending**

The most striking finding from Table 5 is this: the overwhelming majority of existing studies employ only legacy, static datasets such as C-MAPSS, rely solely on FedAvg, and lack XAI integration. No integrated system design in the literature unites federated learning, LEO satellite communication latency management, asynchronous model synchronisation, Explainable AI, and Cross-Fleet “Blind Spot” transfer within a single framework. The Federated Fleet Learning architecture conceptualised in this review targets precisely this holistic gap.

15. Literature Synthesis, Research Gaps, and Future Directions

This comprehensive literature review has synthesised the Predictive Maintenance concept for civil-aviation engines along the axes of federated learning’s data-privacy guarantees, LEO satellites’ communication infrastructure, and Explainable AI’s transparency demands. The high prediction accuracies achieved by 1DCNN-BiLSTM and Transformer architectures can be successfully adapted to airlines’ heterogeneous (Non-IID) data distributions through robust algorithms such as FedProx.

The existing literature, however, retains several research gaps of critical importance that constrain the transition to operational deployment.

The most conspicuous lacuna is the absence of a federated, asynchronous, and space-based integrated framework. To date, federated learning, LEO satellite communication latencies, and Predictive Maintenance have generally been studied in isolation. A clear gap persists in the form of an integrated system design that brings all these components together.

The second critical gap pertains to the empirical validation of sectoral “Blind Spots.” More comprehensive simulations are needed to demonstrate, on datasets such as N-CMAPSS, that airline fleets with entirely disjoint failure types—compressor faults versus turbine degradation—can learn from one another without sharing raw data through Cross-Fleet Knowledge Transfer processes.

A third issue concerns the critical trade-off among privacy, accuracy, and transparency. Differential Privacy (DP) protects the system yet degrades engine-RUL prediction accuracy. Conversely, SHAP-based explanations (XAI) render decisions transparent while potentially leaking information about training data (Alshkeili et al., 2025). How the optimal balance among these three metrics should be calibrated remains a substantial research problem.

A fourth gap lies in the scarcity of hybrid federated algorithm combinations. Studies that combine powerful spatiotemporal architectures such as 1DCNN-BiLSTM with heterogeneity-robust optimisers like FedProx and test feature extraction on Aircraft Edge Units (edge computing) are rare (Arunan et al., 2023).

Finally, the FAA and EASA regulatory wall and certification mismatch cannot be overlooked. Conventional DO-178C standards and EASA AI Roadmap 2.0 demand deterministic behaviour. How complex algorithms distilled from federated networks can be integrated into civil aviation's "Human-in-the-Loop" safety principles with legal compliance remains an open question.

Future research ought to concentrate on a holistic Federated Fleet Learning system that stabilises 1DCNN-BiLSTM architectures with FedProx, operates asynchronously over a Dual LEO Provider network, and meets XAI certification criteria through SHAP analyses. This framework—one that uncompromisingly preserves data privacy, generates collective intelligence, and renders transparent accountability to regulatory authorities—will be determinative in realising aviation's zero-fault vision. The conceptual architecture that this review puts forward provides a robust theoretical foundation for future empirical studies and opens the door to a new collaboration paradigm in the aviation sector.

References

Abdallah, S., & Psomas, C. (2021). The role of satellite communications in 6G: Architecture and key technologies. IEEE Vehicular Technology Magazine, 16(3), 32-41. https://doi.org/10.1109/MVT.2021.3081432

Abdulrahman, S., Tout, H., Ould-Slimane, H., Mourad, A., Talhi, C., & Guizani, M. (2021). A survey on federated learning: The journey from centralized to distributed on-site learning and beyond. IEEE Internet of Things Journal, 8(7), 5476-5497.

Achouch, M., Dimitrova, M., Ziane, K., Sattarpanah Karganroudi, S., Dhouib, R., Ibrahim, H., & Adda, M. (2023). Predictive maintenance and fault monitoring enabled by machine learning: Experimental analysis of a TA-48 multistage centrifugal plant compressor. Applied Sciences, 13(3), 1790. https://doi.org/10.3390/app13031790

Ahn, J., Lee, Y., Kim, N., Park, C., & Jeong, J. (2023). Federated learning for predictive maintenance and anomaly detection using time series data distribution shifts in manufacturing processes. Sensors, 23(17), 7331. https://doi.org/10.3390/s23177331

Ahmed, I., Ahmad, M., Chehri, A., & Jeon, G. (2023). A smart-anomaly-detection system for industrial machines based on feature autoencoder and deep learning. Micromachines, 14(1), 154. https://doi.org/10.3390/mi14010154

Ajay, D. S., Krishnna, S., & Jhajharia, K. (2023). Predictive maintenance of NASA turbofan engines using traditional and ensemble machine learning. In Lecture Notes in Networks and Systems (Vol. 660, pp. 365-376). Springer. https://doi.org/10.1007/978-981-99-1203-2_31

Al Hasib, A., Rahman, A., Khabir, M., & Shawon, M. T. R. (2023). An interpretable systematic review of machine learning models for predictive maintenance of aircraft engine. arXiv preprint arXiv:2309.13310.

Alomari, Y., Andó, M., & Baptista, M. L. (2023). Advancing aircraft engine RUL predictions: An interpretable integrated approach of feature engineering and aggregated feature importance. Scientific Reports, 13, 13466. https://doi.org/10.1038/s41598-023-40315-1

Alshkeili, H. M. H. A., Almheiri, S. J., & Khan, M. A. (2025). Privacy-preserving interpretability: An explainable federated learning model for predictive maintenance in sustainable manufacturing and industry 4.0. AI, 6(6), 117.

Amato, A., & Branco, D. (2025). SemFedXAI: A semantic framework for explainable federated learning in healthcare. Information, 16(6), 435.

Arias Chao, M., Kulkarni, C., Goebel, K., & Fink, O. (2021). Aircraft engine run-to-failure dataset under real flight conditions for prognostics and diagnostics. Data, 6(1), 5. https://doi.org/10.3390/data6010005

Arunan, A., Qin, Y., Li, X., & Yuen, C. (2023). A federated learning-based industrial health prognostics for heterogeneous edge devices using matched feature extraction. IEEE Transactions on Automation Science and Engineering, 21(3), 3065-3079.

Badea, V. E., Zamfiroiu, A., & Boncea, R. (2018). Big data in the aerospace industry. Informatics in Economy, 22(1), 17-24.

Bao, G., & Guo, P. (2022). Federated learning in cloud-edge collaborative architecture: Key technologies, applications and challenges. Journal of Cloud Computing, 11(1), 94.

Baraheem, S., & Yao, Z. (2022). A survey on differential privacy with machine learning and future outlook. arXiv preprint arXiv:2211.10708.

Barbosa, A. M., Ngo, T. V. N., Jafarigol, E., Trafalis, T. B., & Ojoboh, E. P. (2025). Using federated machine learning in predictive maintenance of jet engines. arXiv preprint arXiv:2502.05321.

Barimah, A. K., Onu, O. P., Niculita, O., Cowell, A., & McGlinchey, D. (2025). Scalable data transformation models for physics-informed neural networks in digital twin-enabled PHM applications. Computers, 14(4), 121. https://doi.org/10.3390/computers14040121

Beltrán, E. T. M., Pérez, M. Q., Sánchez, P. M. S., Bernal, S. L., Bovet, G., Pérez, M. G., ... & Celdrán, A. H. (2023). Decentralized federated learning: Fundamentals, state of the art, frameworks, trends, and challenges. IEEE Communications Surveys & Tutorials, 25(4), 2983-3013.

Bemani, M., & Björsell, N. (2023). Federated learning for collaborative predictive maintenance in aviation: A framework. IEEE Access, 11, 14234-14245. https://doi.org/10.1109/ACCESS.2023.3243912

Bharti, S., & McGibney, A. (2021). Privacy-aware resource sharing in cross-device federated model training for collaborative predictive maintenance. IEEE Access, 9, 120367-120379.

Bisanti, G. M., Mainetti, L., Montanaro, T., Patrono, L., & Sergi, I. (2024). Digital twins for aircraft maintenance and operation: A systematic literature review and an IoT-enabled modular architecture. Internet of Things, 24, 100991. https://doi.org/10.1016/j.iot.2023.100991

Biswas, S., Kapoor, P., Chen, J., & Tao, X. (2023). Ensemble neural networks for remaining useful life prediction with uncertainty. arXiv preprint arXiv:2309.12445.

Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., ... & Seth, K. (2017). Practical secure aggregation for privacy-preserving machine learning. In Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security (pp. 1175-1191). ACM.

Chen, J., Jing, H., Chang, Y., & Liu, Q. (2023c). Transfer learning algorithms for bearing remaining useful life prediction: A comprehensive review. Mechanical Systems and Signal Processing, 193, 110239. https://doi.org/10.1016/j.ymssp.2023.110239

Chen, Y., Peng, G., Zhu, Z., & Li, S. (2020). A novel deep learning method based on attention mechanism for bearing remaining useful life prediction. Applied Soft Computing, 86, 105919. https://doi.org/10.1016/j.asoc.2019.105919

Corcuera Bárcena, J. L., Ducange, P., Marcelloni, F., & Renda, A. (2025). Increasing trust in AI through privacy preservation and model explainability: Federated learning of fuzzy regression trees. Information Fusion, 113, 102598. https://doi.org/10.1016/j.inffus.2024.102598

Cummins, L., Sommers, A., Ramezani, S., Mittal, S., Jabour, J. E., Seale, M., & Rahimi, S. (2024). Explainable predictive maintenance: A survey of current methods, challenges and opportunities. IEEE Access, 12, 57574-57602. https://doi.org/10.1109/ACCESS.2024.3351741

Demirbilek, A., Öz, S., & Fidan, Y. (2018). Lojistik Performans Endeksi ve havayolu kargo taşımacılığı ilişkisi: 2007-2016 Türkiye örneği. Ekonomi İşletme ve Yönetim Dergisi, 2(2), 1-17.

Deng, S., & Zhou, J. (2024). Prediction of remaining useful life of aero-engines based on CNN-LSTM-Attention. International Journal of Computational Intelligence Systems, 17(1), 232. https://doi.org/10.1007/s44196-024-00639-w

Ding, Y., Jia, M., Zhuang, J., & Cao, Y. (2022). Transfer learning for remaining useful life prediction across operating conditions based on multisource domain adaptation. IEEE/ASME Transactions on Mechatronics, 27(3), 1539-1551. https://doi.org/10.1109/TMECH.2021.3094744

Dragomir, O. E., Gouriveau, R., Dragomir, F., Minca, E., & Zerhouni, N. (2009, August). Review of prognostic problem in condition-based maintenance. In 2009 European Control Conference (ECC) (pp. 1587-1592). IEEE.

EASA. (2020). EASA AI Roadmap 1.0. European Union Aviation Safety Agency.

EASA. (2023). Artificial intelligence roadmap 2.0: A human-centric approach to AI in aviation. Cologne, Germany: European Union Aviation Safety Agency.

El Ouadrhiri, A., & Abdelhadi, A. (2022). Differential privacy for deep and federated learning: A survey. IEEE Access, 10, 22359-22380. https://doi.org/10.1109/ACCESS.2022.3151670

Eliaz, A., Bouwer, S., & Verhagen, W. J. C. (2022). A systematic literature review of predictive maintenance for defence fixed-wing aircraft sustainment and operations. Sensors, 22(18), 7070. https://doi.org/10.3390/s22187070

El-Rashidy, N., Derhem, A., Abuhmed, T., Elbaz, O., El-Bakry, H. M., & El-Fishawy, N. (2022). Remaining useful life prediction of an aircraft turbofan engine using deep layer recurrent neural networks. Actuators, 11(3), 67. https://doi.org/10.3390/act11030067

Eutelsat Communications. (2023). Eutelsat completes combination with OneWeb. https://www.eutelsat.com/en/newsroom/press-releases/eutelsat-completes-combination-with-oneweb.html

FAA. (2011). Advisory Circular AC 20-115C: Airborne Software Assurance. Federal Aviation Administration.

Fan, F., Xiong, J., Li, M., & Wang, G. (2021). On interpretability of artificial neural networks: A survey. IEEE Transactions on Radiation and Plasma Medical Sciences, 5(6), 741-760.

Fernandes, M., Corchado, J. M., & Marreiros, G. (2022). Machine learning techniques applied to mechanical fault diagnosis and fault prognosis in the context of real industrial manufacturing use-cases: A systematic literature review. Applied Intelligence, 52(12), 14246-14280. https://doi.org/10.1007/s10489-022-03344-3

Fu, S., & Avdelidis, N. P. (2023). Prognostic and health management of critical aircraft systems and components: An overview. Sensors, 23(19), 8124. https://doi.org/10.3390/s23198124

Geng, J., Mou, Y., Jin, H., Wang, X., Luo, M., & Yuan, C. (2023). Improved gradient inversion attacks and defenses in federated learning. IEEE Transactions on Big Data, 9(5), 1337-1350. https://doi.org/10.1109/TBDATA.2023.3241254

Guo, J., Lei, S., & Du, B. (2024). MHT: A multiscale hourglass-transformer for remaining useful life prediction of aircraft engine. Engineering Applications of Artificial Intelligence, 128, 107519. https://doi.org/10.1016/j.engappai.2023.107519

Guo, P., Zeng, S., Chen, W., Zhang, X., Ren, W., Zhou, Y., & Qu, L. (2025). A new federated learning framework against gradient inversion attacks. In Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 39, No. 16, pp. 16969-16977).

Hakami, A. (2024). Strategies for overcoming data scarcity, imbalance, and feature selection challenges in machine learning models for predictive maintenance. Scientific Reports, 14(1), 9645.

Hızıroğlu, A., Pişirgen, A., Özcan, M., & İlter, H. K. (2022). Artificial intelligence in healthcare industry: A transformation from model-driven to knowledge-driven decision support systems. Artificial Intelligence Theory and Applications, 2(1), 41-58.

Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. Neural Computation, 9(8), 1735-1780.

Hong, C. W., Lee, C., Lee, K., Ko, M. S., Kim, D. E., & Hur, K. (2020). Remaining useful life prognosis for turbofan engine using explainable deep neural networks with dimensionality reduction. Sensors, 20(22), 6626. https://doi.org/10.3390/s20226626

Huang, C. G., Huang, H. Z., Li, Y. F., & Peng, W. (2023). A health state-related ensemble deep learning method for aircraft engine remaining useful life prediction. Applied Soft Computing, 135, 109925. https://doi.org/10.1016/j.asoc.2023.109925

Huang, Y., Gupta, S., Song, Z., Li, K., & Arber, S. (2021). Evaluating gradient inversion attacks and defenses in federated learning. In Advances in Neural Information Processing Systems (Vol. 34, pp. 7232-7241).

IATA. (2024). Airline maintenance cost executive commentary: FY2023 MCX report. International Air Transport Association. https://www.iata.org/mcx

Jagatheesaperumal, S. K., Rahouti, M., Ahmad, K., Al-Fuqaha, A., & Guizani, M. (2024). Enabling trustworthy federated learning in industrial IoT: Bridging the gap between interpretability and robustness. IEEE Internet of Things Magazine, 7(4), 108-114. https://doi.org/10.1109/IOTM.001.2400003

Jakubowski, J., Stanisz, P., Bobek, S., & Nalepa, G. J. (2022). Performance of explainable AI methods in asset failure prediction. In Lecture Notes in Computer Science (Vol. 13353, pp. 472-485). Springer. https://doi.org/10.1007/978-3-031-08760-8_40

Kabashkin, I. (2024). Integration of foundation models and federated learning in AIoT-based aircraft health monitoring systems. Mathematics, 12(21), 3428. https://doi.org/10.3390/math12213428

Kabashkin, I., Perekrestov, V., & Pivovar, M. (2025). AI-driven fault detection and maintenance optimization for aviation technical support systems. Processes, 13(3), 666. https://doi.org/10.3390/pr13030666

Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., ... & Zhao, S. (2021). Advances and open problems in federated learning. Foundations and Trends in Machine Learning, 14(1-2), 1-210. https://doi.org/10.1561/2200000083

Kamei, S., & Taghipour, S. (2023). A comparison study of centralized and decentralized federated learning approaches utilizing the transformer architecture for estimating remaining useful life. Reliability Engineering and System Safety, 233, 109130. https://doi.org/10.1016/j.ress.2023.109130

Karami, A., Kassab, R., & Sezgin, A. (2023). Low-latency collaborative predictive maintenance: Over-the-air federated learning in noisy industrial environments. Sensors, 23(18), 7840. https://doi.org/10.3390/s23187840

Karimireddy, S. P., Kale, S., Mohri, M., Reddi, S. J., Stich, S. U., & Suresh, A. T. (2020). SCAFFOLD: Stochastic controlled averaging for federated learning. In Proceedings of the 37th International Conference on Machine Learning (pp. 5132-5143). PMLR.

Kim, S., Choi, J. H., & Kim, N. H. (2022a). Data-driven prognostics with low-fidelity physical information for digital twin: Physics-informed neural network. Structural and Multidisciplinary Optimization, 65, 255. https://doi.org/10.1007/s00158-022-03348-0

Konečný, J., McMahan, H. B., Yu, F. X., Richtárik, P., Suresh, A. T., & Bacon, D. (2016). Federated learning: Strategies for improving communication efficiency. arXiv preprint arXiv:1610.05492.

Kumru, S., & Kasımoğlu, M. (2022). İşletmelerde Dijital Dönüşümün Stratejik Yönetimi: Bir Alan Araştırması. Bilgi Ekonomisi ve Yönetimi Dergisi, 17(2), 139-159.

Kusiak, A. (2024). Federated explainable artificial intelligence (fXAI): A digital manufacturing perspective. International Journal of Production Research, 62(1-2), 171-182. https://doi.org/10.1080/00207543.2023.2238083

Landau, D., de Pater, I., Mitici, M., & Saurabh, N. (2026). Federated learning framework for collaborative remaining useful life prognostics: An aircraft engine case study. Future Generation Computer Systems, 174, 107945.

Lei, Y., Yang, B., Jiang, X., Jia, F., Li, N., & Nandi, A. K. (2021). Challenges and opportunities of AI-enabled monitoring, diagnosis & prognosis: A review. Chinese Journal of Mechanical Engineering, 33, 56. https://doi.org/10.1186/s10033-021-00570-7

Li, C., Li, S., Feng, Y., Gryllias, K., Gu, F., & Pecht, M. (2024a). Small data challenges for intelligent prognostics and health management: A review. Artificial Intelligence Review, 57(8), 214.

Li, F., Zhang, L., Chen, B., Gao, D., Cheng, Y., & Zhang, Y. (2018). A light gradient boosting machine for remaining useful life estimation of aircraft engines. In 2018 IEEE Intelligent Transportation Systems Conference (ITSC) (pp. 3562-3567). IEEE. https://doi.org/10.1109/ITSC.2018.8569801

Li, L., Fan, Y., Tse, M., & Lin, K. (2020a). A review of applications in federated learning. Computers & Industrial Engineering, 149, 106854.

Li, S., Jin, N., Dogani, A., Yang, Y., Zhang, M., & Gu, X. (2024b). Enhancing LightGBM for industrial fault warning: An innovative hybrid algorithm. Processes, 12(1), 221.

Li, T., Sahu, A. K., Zaheer, M., Sanjabi, M., Talwalkar, A., & Smith, V. (2020b). Federated optimization in heterogeneous networks. In Proceedings of Machine Learning and Systems (Vol. 2, pp. 429–450).

Li, X., Zhang, W., Ding, Q., & Sun, J. Q. (2022b). Domain adaptive remaining useful life prediction with transformer. IEEE Transactions on Instrumentation and Measurement, 71, 1–13. https://doi.org/10.1109/TIM.2022.3163156

Liu, H., Chen, G., & Wang, X. (2022). A fair and efficient hybrid federated learning framework based on XGBoost for distributed power prediction. arXiv preprint arXiv:2201.02783.

Liu, Y., Zhang, L., Nie, J., & Wang, Z. (2024). Federated learning-oriented edge computing framework for the IIoT. Sensors, 24(13), 4182. https://doi.org/10.3390/s24134182

Llasag Rosero, R., Silva, C., Ribeiro, B., Albisser, M., Brutsche, M., & Arias Chao, M. (2025). Label synchronization strategies for hybrid federated learning. Reliability Engineering and System Safety, 256, 110751. https://doi.org/10.1016/j.ress.2024.110751

Lopez-Ramos, L. M., Garcia-Martinez, A., & Ruiz-Gonzalez, R. (2024). Interplay between federated learning and explainable artificial intelligence: A scoping review. arXiv preprint arXiv:2411.05874.

Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. In Advances in Neural Information Processing Systems (Vol. 30, pp. 4765–4774).

Ma, Q., Zhang, L., Chen, Y., & Liu, X. (2021). Remaining useful life estimation of aircraft engines based on deep convolutional neural network and LightGBM combination model. International Journal of Computational Intelligence Systems, 14, 20. https://doi.org/10.1007/s44196-021-00020-1

McMahan, B., Moore, E., Ramage, D., Hampson, S., & Arcas, B. A. (2017). Communication-efficient learning of deep networks from decentralized data. In Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS) (pp. 1273–1282). PMLR.

Michel, F., Trevisan, M., Giordano, D., & Bonaventure, O. (2022). A first look at starlink performance. In Proceedings of the 22nd ACM Internet Measurement Conference (pp. 130–136).

Milasheuski, U., Baraldi, P., Zio, E., & Savazzi, S. (2024). Federated generative models for predictive maintenance in industrial environments. In 8th International Conference on System Reliability and Safety (ICSRS) (pp. 156–161). IEEE.

Mobley, R. K. (2002). An introduction to predictive maintenance. Elsevier.

Mo, H., Ijaz, H., & Wang, X. (2023). A hybrid deep learning approach for RUL prediction of turbofan engines using 1D-CNN and Bi-LSTM. Journal of Mechanical Science and Technology, 37(4), 1899–1910. https://doi.org/10.1007/s12206-023-0328-7

Mohan, N., Ferguson, A. E., Cech, H., Bose, R., Renatin, P. R., Marina, M. K., & Ott, J. (2024). A multifaceted look at starlink performance. In Proceedings of the ACM Web Conference 2024 (pp. 2723-2734).

Nemat Saberi, A., Sedighi, M., & Asghari, V. (2022). LightGBM-based fault diagnosis of rotating machinery under non-stationary conditions. IEEE Access, 10, 81910-81925. https://doi.org/10.1109/ACCESS.2022.3196167

Nguyen, D. C., Ding, M., Pathirana, P. N., Seneviratne, A., Li, J., & Poor, H. V. (2021). Federated learning for industrial internet of things in future industries. IEEE Internet of Things Journal, 8(21), 16307-16328. https://doi.org/10.1109/JIOT.2021.3091652

Nguyen, Q. M., Le, N. K., & Nguyen, L. M. (2023, June). Scalable and secure federated xgboost. In ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) (pp. 1-5). IEEE.

Oliver Wyman. (2025). Global fleet and MRO market forecast 2025-2035. https://www.oliverwyman.com/our-expertise/insights/2025/feb/global-fleet-and-mro-market-forecast-2025-2035.html

Pan, D., Liu, Z., & Zhou, J. (2023). Interpretable fault diagnosis of aero-engine sensors based on SHAP values. Measurement Science and Technology, 34(5), 055108. https://doi.org/10.1088/1361-6501/acbf14

Pan, S. J., & Yang, Q. (2010). A survey on transfer learning. IEEE Transactions on Knowledge and Data Engineering, 22(10), 1345-1359.

Park, P., Di Marco, P., Shin, H., & Bang, J. (2019). Fault detection and diagnosis using combined autoencoder and long short-term memory network. Sensors, 19(21), 4612. https://doi.org/10.3390/s19214612

Qiu, S., Cui, X., Ping, Z., Shan, N., Li, Z., Bao, X., & Xu, X. (2023). Deep learning techniques in intelligent fault diagnosis and prognosis for industrial systems: A review. Sensors, 23(3), 1305. https://doi.org/10.3390/s23031305

Öz, S., Dindarik N., Duman, F. (2019). Schumpeteryan Girişimcilik ve 21. Yüzyıl Girişimcilik Yaklaşımları: Y Kuşağı Girişimcileri Üzerine Nitel Bir Araştırma, opus (13), pp. 1059-1091.

Öz, S. (2020). Genel Teknoloji ve Dijital Dönüşüm Teorisine Giriş. Book Chapter in "Teknolojik ve Dijital Dönüşüm", Nobel, pp. 169-192.

Reddi, S. J., Charles, Z., Zaheer, M., Garrett, Z., Rush, K., Konečný, J., Kumar, S., & McMahan, H. B. (2021). Adaptive federated optimization. In Proceedings of the 9th International Conference on Learning Representations (ICLR).

Salih, A., Boscoe, F., & Raj, R. (2023). A perspective on explainable artificial intelligence methods: SHAP and LIME. arXiv preprint arXiv:2305.02012.

Scarselli, G., & Nicassio, F. (2025). Machine learning for structural health monitoring of aerospace structures: A review. Sensors, 25(19), 6136.

Serradilla, O., Zugasti, E., Rodriguez, J., & Zurutuza, U. (2022). Deep learning models for predictive maintenance: A survey, comparison, challenges and prospects. Applied Intelligence, 52(10), 10934–10964.

SES S.A. (2024). SES launches SES Open Orbits™ inflight connectivity network. https://www.ses.com/press-release/ses-launches-ses-open-orbitstm-inflight-connectivity-network

Sherifi, A. (2024). Turbofan engine remaining useful life prediction based on bidirectional LSTM. arXiv preprint arXiv:2411.16422.

Son, T., Hong, D., & Kim, B. (2023). Multi-output classification based on convolutional neural network model for untrained compound fault diagnosis of rotating machinery. Sensors, 23(6), 3153. https://doi.org/10.3390/s23063153

SpaceX. (2024). Inter-satellite laser links deployment. https://www.spacex.com/updates/

SpaceX. (2025). Starlink constellation status. https://www.starlink.com/map

Stanton, I., Munir, K., Ikram, A., & El-Bakry, M. (2023). Predictive maintenance analytics and implementation for aircraft: Challenges and opportunities. Systems Engineering, 26(2), 216–237. https://doi.org/10.1002/sys.21651

Starlink. (2024). Network performance update. https://starlink.com/updates

Theissler, A., Spinnato, F., Schlegel, U., & Guidotti, R. (2022). Explainable AI for time series classification: A review, taxonomy and research directions. IEEE Access, 10, 100700–100724.

Tian, Q., & Wang, H. (2020). An ensemble learning and RUL prediction method based on bearings degradation indicator construction. Applied Sciences, 10(1), 346.

Uçar, A., Karaköse, M., & Kırımça, N. (2024). Artificial intelligence for predictive maintenance applications: Key components, trustworthiness, and future trends. Applied Sciences, 14(2), 898. https://doi.org/10.3390/app14020898

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. In Advances in Neural Information Processing Systems (Vol. 30, pp. 5998–6008).

Velasco-Loera, F., Alcaraz-Mejia, M., & Chavez-Hurtado, J. L. (2025). An interpretable hybrid fault prediction framework using XGBoost and Bayesian networks. Applied Sciences, 15(18), 10164. https://doi.org/10.3390/app151810164

Verhagen, W. J. C., de Boer, L. W. M., & Curran, R. (2023). Condition-based maintenance in aviation: Challenges and opportunities. Aerospace, 10(9), 762. https://doi.org/10.3390/aerospace10090762

Von Wahl, L., Heidenreich, N., Mitra, P., Nolting, M., & Tempelmeier, N. (2024). Data disparity and temporal unavailability aware asynchronous federated learning for predictive maintenance on transportation fleets. Proceedings of the AAAI Conference on Artificial Intelligence, 15420–15428.

Wahid, A., Breslin, J. G., & Intizar, M. A. (2022). Prediction of machine failure in Industry 4.0: A hybrid CNN-LSTM framework. Applied Sciences, 12(9), 4221. https://doi.org/10.3390/app12094221

Wang, J., Liu, Q., Liang, H., Joshi, G., & Poor, H. V. (2020). Tackling the objective inconsistency problem in heterogeneous federated optimization. Advances in Neural Information Processing Systems, 33, 7611-7623.

Wen, P., Ye, Z., Li, Y., Chen, S., & Xie, Y. (2023). Physics-informed neural networks for prognostics and health management of lithium-ion batteries. IEEE Transactions on Neural Networks and Learning Systems, 34(12), 10654-10668. https://doi.org/10.1109/TNNLS.2022.3167960

Wen, Q., Zhou, T., Zhang, C., Chen, W., Ma, Z., Yan, J., & Sun, L. (2022). Transformers in time series: A survey. arXiv preprint arXiv:2202.07125.

Xia, G., Chen, J., Yu, C., & Ma, J. (2023). Poisoning attacks in federated learning: A survey. IEEE Access, 11, 10708-10722. https://doi.org/10.1109/ACCESS.2023.3238823

Xia, Y., Yu, W., & Li, Q. (2025). Byzantine-resilient federated learning via distributed optimization. In 2025 33rd European Signal Processing Conference (EUSIPCO) (pp. 2672-2676). IEEE.

Xie, L., Liu, J., Lu, S., Chang, T. H., & Shi, Q. (2022). An efficient learning framework for federated XGBoost using secret sharing and distributed optimization. ACM Transactions on Intelligent Systems and Technology (TIST), 13(5), 1-28.

Xiong, M., Wang, H., Fu, Q., & Xu, Y. (2021). Digital twin-driven aero-engine intelligent predictive maintenance. The International Journal of Advanced Manufacturing Technology, 114(11), 3751-3761.

Xu, H., Zhang, Z., Di, S., Liu, B., Alharthi, K. A., & Cao, J. (2024). Fedfa: A fully asynchronous training paradigm for federated learning. arXiv preprint arXiv:2404.11015.

Xu, J., Zhang, Z., & Hu, R. (2025). Achieving byzantine-resilient federated learning via layer-adaptive sparsified model aggregation. In 2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) (pp. 1508-1517). IEEE.

Xu, Q., Wu, M., Khoo, E., Chen, Z., & Li, X. (2023b). A hybrid ensemble deep learning approach for early prediction of battery remaining useful life. IEEE/CAA Journal of Automatica Sinica, 10(1), 177-187.

Yang, L., Tham, C. K., & Guo, S. (2023). FedPreM: A novel federated reinforcement learning framework for predictive maintenance. In GLOBECOM 2023-2023 IEEE Global Communications Conference (pp. 7556-7561). IEEE.

Yıldız, T. K., Yurtay, N., & Öneç, B. (2021). Classifying anemia types using artificial learning methods. Engineering Science and Technology, an International Journal, 24(1), 50-70.

Yu, B., Guo, H., & Shi, J. (2025). Remaining useful life prediction based on hybrid CNN-BiLSTM model with dual attention mechanism. International Journal of Electrical Power & Energy Systems, 172, 111152.

Zhang, B., Zheng, K., Huang, Q., Feng, S., Zhou, S., & Zhang, Y. (2020). Aircraft engine prognostics based on informative sensor selection and adaptive degradation modeling with functional principal component analysis. Sensors, 20(3), 920.

Zhang, W., & Li, X. (2022). Data privacy preserving federated transfer learning in machinery fault diagnostics using prior distributions. Structural Health Monitoring, 21(4), 1329–1344.

Zhang, W., Yang, D., & Wang, H. (2019). Data-driven methods for predictive maintenance of industrial equipment: A survey. IEEE Systems Journal, 13(3), 2213–2227.

Zhao, R., Yan, R., Chen, Z., Mao, K., Wang, P., & Gao, R. X. (2019). Deep learning and its applications to machine health monitoring. Mechanical Systems and Signal Processing, 115, 213–237.

Zhou, F., Yang, Y., Wang, C., & Hu, X. (2023). Federated learning based fault diagnosis driven by intra-client imbalance degree. Entropy, 25(4), 606.

Zuo, S., Yan, X., Fan, R., Hu, H., Shan, H., Quek, T. Q., & Zhao, P. (2025). Federated learning resilient to byzantine attacks and data heterogeneity. IEEE Transactions on Mobile Computing.