International Journal of Modern Research in Engineering and Technology (IJMRET)
www.ijmret.org Volume 10 Issue 10 ǁ October 2025.

"Exploring Machine Learning Approaches for Fault Prediction in Airbus A320 Air Conditioning Systems: A Qualitative Review"

Manyanga David Victor Vanessa[1] Wu Honglan[2] Magige Medard Makile
[1,2,3]*(Dept. of Civil Aviation*
Nanjing University of Aeronautics and Astronautics, Nanjing 210016, China)

ABSTRACT: *The aviation industry is transitioning from reactive to predictive maintenance as aircraft subsystems grow increasingly complex. This paper presents a qualitative review of machine learning (ML) approaches for fault prediction in the Airbus A320 Air Conditioning System (ACS), emphasizing interpretability, socio-technical integration, and regulatory alignment. Unlike prior research that prioritizes algorithmic performance metrics, this study consolidates existing literature and industry practices to propose a structured framework for integrating interpretable ML into predictive maintenance workflows.The analysis highlights the strengths of Random Forests and Decision Trees in handling multidimensional sensor data while maintaining transparency. Supported by SHAP-based explanations, these models can provide actionable insights that build technician trust and facilitate regulatory acceptance. Beyond technical performance, the framework addresses organizational readiness, workflow compatibility, and compliance with aviation standards, including DO-178C, ARP4754A, and ICAO's Safety Management System (SMS).A conceptual validation pathway is proposed, progressing from retrospective analysis to shadow-mode deployment and monitored roll-out, ensuring human-in-the-loop oversight throughout. The study concludes that interpretable ML, when combined with regulatory governance and human-centered design, offers a viable pathway toward safer, cost-efficient, and operationally sustainable maintenance in commercial aviation.*

KEYWORDS -*Airbus A320, Air Conditioning System (ACS), Predictive Maintenance, Machine Learning, Explainable AI, Aviation Safety.*

I. INTRODUCTION

The aviation industry has always prioritized safety, reliability, and cost-efficiency. With the increasing complexity and interconnection of modern aircraft systems, traditional maintenance practices centered on scheduled inspections or reactive diagnostics have shown significant limitations. Particularly for critical subsystems such as the Air Conditioning System (ACS) in the Airbus A320, unanticipated failures not only compromise passenger comfort and safety but can also cause delays, operational disruptions, and increased maintenance costs. As global air traffic continues to rise, operators face growing pressure to maximize fleet availability while minimizing downtime, necessitating a shift toward more intelligent, adaptive maintenance strategies.

In this evolving landscape, machine learning (ML) has gained attention as a tool for predictive maintenance. Unlike rule-based or threshold-driven diagnostics, ML algorithms can learn patterns from complex, high-volume datasets such as real-time sensor streams, flight records, and historical maintenance logs. By jointly analyzing temporal, environmental, and mechanical variables, ML methods can anticipate potential failures before they manifest, thereby reducing unscheduled interventions and optimizing resource allocation. For example, Bhola and Malhotra (2020)[1] showed that predictive paradigms can significantly improve operational efficiency, while Jennions et al. (2022)[2] and Zhang et al. (2024)[3] highlighted the potential of ML models to address rare fault detection in highly imbalanced datasets.

The Airbus A320, one of the most widely used commercial aircraft worldwide, provides a compelling test-bed for predictive maintenance research. Its ACS integrates multiple sensors that record airflow, temperature, compressor activity, and valve positions, generating multidimensional data during every flight. Manually analyzing these signals to detect early indicators of degradation is both labor-intensive and error-prone. Supervised learning techniques such as Random Forests, Decision Trees, and Logistic Regression have therefore been investigated for their ability to capture subtle correlations among system variables. When embedded into Aircraft Health Monitoring Systems (AHMS) and linked with digital twin frameworks, such models support condition-based maintenance by enabling proactive interventions without disrupting operations.

Despite these advances, important challenges remain. Most prior studies focus narrowly on algorithmic accuracy, often neglecting socio-technical factors such as technician usability, workflow integration, and compliance with regulatory standards like DO-178C and Safety Management Systems (SMS). As a result, the pathway from research prototypes to operational deployment remains underexplored.

This paper addresses that gap by presenting a qualitative synthesis of existing work and developing a conceptual framework for the safe and interpretable integration of ML-based fault prediction into Airbus A320 ACS maintenance workflows. Rather than reporting new model training experiments, we consolidate prior findings to evaluate feasibility, interpretability, and practical barriers to adoption. The framework emphasizes the convergence of technical model design, human-centered usability, and regulatory assurance, contributing a holistic perspective often absent from predictive maintenance literature.

Contributions

This paper makes the following contributions to the literature on predictive maintenance in aviation:

Qualitative Synthesis of ML Approaches — consolidates prior studies on machine learning-based fault prediction with a focus on the Airbus A320 Air Conditioning System (ACS), highlighting patterns, challenges, and open gaps.

Socio-Technical Integration Framework — proposes a conceptual pathway for embedding interpretable ML into ACS maintenance workflows, emphasizing technician usability, workflow compatibility, and stakeholder trust.

Regulatory Alignment Perspective — links predictive maintenance practices with safety assurance frameworks such as DO-178C and Safety Management Systems (SMS), offering insights into certification and compliance challenges.

Future Research and Practical Directions — identifies limitations in current approaches, such as imbalanced datasets and limited real-world validation, and provides recommendations to guide both academic research and industrial adoption.

II. LITERATURE REVIEW

2.1 Technical Advances in ML for Predictive Maintenance

Machine learning has been increasingly applied to predictive maintenance in both HVAC and aviation domains. Arshad, Tyagi, and Kalia (2023) [4] demonstrated that ensemble methods such as Random Forests perform robustly on noisy, multidimensional datasets, while Bhola and Malhotra (2020) [1] showed that supervised models can detect anomalies earlier than rule-based systems in air conditioning maintenance. Similar trends are evident in aviation-specific contexts, where supervised models have been applied to sensor-rich subsystems like the A320's Air Conditioning System (ACS). These approaches highlight ML's capacity to process high-frequency flight data and detect subtle precursors of failure. However, most of these contributions prioritize numerical performance metrics (e.g., accuracy, F1-scores) over issues of explainability or deployment readiness.

2.2 Data Challenges in Aviation Fault Prediction

One persistent challenge is rare fault detection in imbalanced datasets. Jennions, King, and Skaf (2022) and Dangut et al. (2023) [2] emphasized that most recorded flight data represent healthy system states, limiting model generalization to rare but critical failures. To address this, hybrid methods and data fusion strategies have been proposed, integrating sensor readings with historical maintenance records. Zhang, Li, and Wang (2024) [3] further advocate for collaborative data-sharing ecosystems to enhance scalability and improve diagnostic accuracy across fleets. These contributions underscore the necessity of unified data infrastructures but leave open questions around data governance and standardization.

2.3 Explainability and Operator Trust

While predictive performance is important, model interpretability has emerged as a decisive factor in aviation adoption. Amann et al. (2020) and Miller (2019) argue that explainability is central to operator trust, particularly in safety-critical settings. Recent studies (Wang et al., 2024; Liu et al., 2025; Raj & Thomas, 2024) demonstrate the use of SHAP (SHapley Additive Explanations) in aerospace and HVAC time-series contexts. These works show how SHAP can identify influential features such as compressor temperature, precooler pressure, and cabin zone readings. Yet, García et al. (2023) and Kim et al. (2023) caution that there remains a gap between mathematically faithful explanations and practical interpretability for technicians and regulators.

2.4 Regulatory and Governance Perspectives

The integration of predictive maintenance tools into aviation must align with regulatory frameworks. Bhola and Malhotra (2020) [1] and Striim (2024) [5] note that compliance with standards such as EASA's AI Roadmap (2020), ICAO's Safety Management System (SMS), and DO-178C requirements presents significant hurdles, given ML's adaptive nature. WeShield (2024) [7] suggests hybrid assurance strategies, where probabilistic predictions are paired with deterministic safety checks to meet certification demands. Beyond model validation, studies emphasize the importance of governance mechanisms such as continuous monitoring, audit trails, and alert triage protocols to ensure accountability in operational contexts.

Collectively, these studies establish a strong technical foundation for ML-based fault prediction in aviation. However, gaps remain in three critical areas: (1) limited adoption of interpretable methods, (2) insufficient attention to human factors and workflow integration, and (3) unresolved tensions between adaptive ML systems and static regulatory requirements. Addressing these gaps provides the rationale for the socio-technical integration framework developed in this paper.

Table 01 : Summary of Reviewed Studies on ML for Predictive Maintenance in Aviation and Related Domains

Author(s), Year	Domain / System	Data Type & Source	ML Methods Used	Interpretability Techniques	Key Findings	Limitations / Gaps
Bhola & Malhotra, 2020 [1]	HVAC / Air Conditioning	Historical maintenance & sensor data	Supervised Learning (Decision Trees, Logistic Regression)	None reported	Early anomaly detection improved maintenance scheduling efficiency compared to conventional methods.	Focused on accuracy; did not address interpretability or regulatory adoption.
Arshad, Tyagi & Kalia, 2025 [4]	HVAC & related systems	Simulated + real sensor data	Ensemble methods (Random Forests, Gradient Boosting)	Feature importance (basic)	Demonstrated robustness of ensemble models to noisy, multi-dimensional data; highlighted predictive potential.	No discussion of human factors, trust, or workflow integration.
Jennions, King & Skaf, 2022; Dangut et al., 2023 [2]	Aviation (general)	Flight logs & operational sensor data	Hybrid ML, Deep Learning	Limited (mostly model outputs)	Proposed methods for rare fault detection; emphasized data imbalance challenges.	Did not evaluate explainability or practical adoption; limited generalizability to ACS-specific faults.
Zhang, Li & Wang, 2024 [3]	Aviation (fleet-level)	Cross-airline operational & maintenance data	Data fusion + ML frameworks	None reported	Advocated for unified data-sharing ecosystems to improve scalability and diagnostic accuracy.	Governance, data privacy, and regulatory integration remain unresolved.
Amann et al., 2020; Miller, 2019	Explainable AI (general, applied to safety-critical)	Conceptual / secondary sources	XAI frameworks (model-agnostic)	SHAP, LIME, transparency principles	Argued that explainability is critical for adoption in regulated domains.	No aviation-specific experiments; mainly conceptual discussions.
Sharma & Mistry, 2023 [6]	Predictive Maintenance (multi-domain)	Case studies, industrial data	Supervised ML	None reported	Showed that ML systems without intuitive interfaces often fail in adoption, even if technically accurate.	Overlooked interpretability; no aviation-specific context.

Author(s), Year	Domain / System	Data Type & Source	ML Methods Used	Interpretability Techniques	Key Findings	Limitations / Gaps
Wellsfield, 2024 [7] (industry whitepaper)	Aviation predictive maintenance	Conceptual / industrial commentary	Hybrid probabilistic – deterministic	Audit trails, safety checks (conceptual)	Suggested blending ML predictions with deterministic safety checks for certification.	Non-peer-reviewed; lacks empirical validation.
Wang et al., 2024; Liu et al., 2025; Raj & Thomas, 2024	Aerospace & HVAC diagnostics	Sensor-based time series	Random Forests, Deep Learning	SHAP	Demonstrated SHAP's ability to identify influential variables (e.g., compressor temp, precooler pressure).	Faithfulness vs interpretability gap persists; explanations not yet technician-friendly.
Garcia et al., 2023; Kim et al., 2023	XAI in regulated settings	Human-in-the-loop studies	Multiple ML models	SHAP, LIME, user studies	Found that operator trust improves with transparent explanations, but usability barriers remain.	Small sample sizes; limited focus on aviation-specific contexts.

The summary in Table 01 highlights both the breadth and fragmentation of current research on ML-based predictive maintenance. Across domains, ensemble and hybrid models have consistently shown promise in handling complex, multidimensional data, while deep learning approaches are increasingly applied to rare fault detection. However, interpretability methods are either absent or treated superficially in most studies, with only a limited number explicitly applying SHAP or similar XAI techniques. Furthermore, while industry reports and conceptual studies acknowledge the importance of regulatory compliance, few peer-reviewed works systematically map predictive models to certification frameworks such as DO-178C or ICAO SMS. Finally, the majority of studies emphasize algorithmic accuracy but neglect socio-technical factors, such as technician usability, workflow integration, and governance mechanisms.

Taken together, these findings reveal a clear research gap: although ML offers strong technical capabilities for predictive maintenance, the lack of integrated, interpretable, and regulation-aligned approaches limits real-world adoption. Addressing this gap is the central motivation for the socio-technical integration framework proposed in the following section.

III. METHODOLOGY

Qualitative and Conceptual Framework

This study adopts a qualitative and conceptual methodology, rather than an empirical modeling approach. No new data collection, algorithm training, or quantitative benchmarking was conducted. Instead, the paper synthesizes existing literature and established industry practices to propose a structured framework for interpretable machine learning (ML) in the predictive maintenance of the Airbus A320 Air Conditioning System (ACS). The methodology is therefore oriented toward framework development, emphasizing data ecosystem characterization, feature relevance, model interpretability, and regulatory alignment.

The objective is not to report experimental performance metrics, but to demonstrate how ML concepts can be operationally integrated into aviation maintenance workflows. Accordingly, the framework is illustrated through conceptual figures and scenarios, designed to clarify socio-technical interactions between data, models, technicians, and compliance structures. These elements establish the basis for later discussion of interpretability, organizational fit, and safety assurance.

1. Data Ecosystem and Operational Context

This study adopts a qualitative, conceptual methodology. No new empirical data collection, model training, or performance benchmarking was undertaken. Instead, the analysis synthesizes existing literature and established industry practices to describe the Airbus A320 Air Conditioning System (ACS) data ecosystem, relevant preprocessing strategies, model selection rationales, and the conceptual application of explainability techniques. The objective is to propose a structured integration framework for interpretable machine learning (ML) within predictive maintenance workflows.

The effective application of ML in aviation maintenance is contingent upon the quality, completeness, and contextual richness of the underlying data. The ACS of the Airbus A320 comprises multiple components such as compressors, valves, ducts, and temperature regulation units monitored through embedded sensors. These sensors capture operational parameters in real time, including airflow rates, compressor temperatures, and engine rotational speeds.

The primary categories of data informing predictive analytics include:

Sensor Data – High-frequency time-series measurements such as PACK_FLOW_R1/R2, PACK_COMPR_T1/T2, discharge air temperatures, and engine RPM values (N11–N22).

Flight Operation Logs – Contextual metadata detailing the operational environment, including altitude, flight phase, Mach number, and ambient temperature.

Maintenance Records – Historical documentation of faults, component replacements, and corrective interventions, providing a longitudinal perspective on system performance.

The integration of these three data domains enables a multi-dimensional representation of system health, supporting the identification of both transient anomalies and recurrent failure patterns. Prior studies (Jennions, King, & Skaf, 2022; Zhang, Li, & Wang, 2024) underscore the diagnostic value of correlating real-time sensor telemetry with maintenance history to enhance fault detection fidelity.

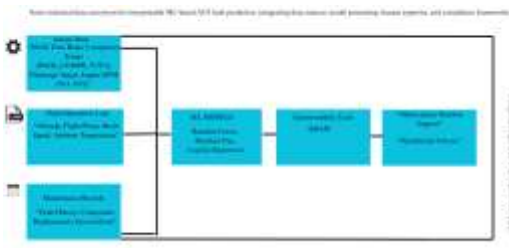

Figure 1. Socio-technical data ecosystem for interpretable ML-based ACS fault prediction, integrating data sources, model processing, human expertise, and compliance frameworks.

Socio-technical data ecosystem for interpretable ML-based fault prediction in the Airbus A320 Air Conditioning System (ACS). The framework integrates sensor data, flight operation logs, and maintenance records through preprocessing, feature engineering, and model processing. Outputs are enhanced by human expertise and governed by compliance frameworks (DO-178C, ARP4754A, SMS) to support operational decision-making.

When combined, these data streams form a multidimensional view of aircraft health, enabling the detection of both transient anomalies and recurring failure patterns. Recent studies (Jennions, King, & Skaf, 2022;[2] Zhang, Li, & Wang, 2024[3]) emphasize the value of cross-referencing sensor telemetry with historical maintenance logs to yield a richer context for diagnostic reasoning.

2. Feature Engineering and Relevance to Technicians

Instead of just aiming for the most precise algorithms, this study sees the preparation of data as a crucial first step in how machine learning interacts with complex aviation data. Sometimes, sensor data is missing because of issues like dropped signals or delays. To handle this, missing values are filled in using methods like averaging or KNN, and maintenance teams also watch out for missing data as a sign of sensor problems.

Transforming the data, such as scaling it or changing its format, helps make the information easier for humans to understand. For example, instead of using complex numbers, engineers prefer inputs like average airflow over 30 seconds, which match their everyday understanding of system behavior.

Further, techniques like Principal Component Analysis (PCA) reduce the number of data features, making analysis easier. More importantly, PCA highlights which system parts are most important in predicting faults like certain temperatures or component pressures. These insights can help engineers focus on specific areas for inspection or design improvements, creating a helpful link between data and actual operations.

Table 1: Key Features for Airbus A320 ACS Fault Prediction – Description, diagnostic relevance, and typical behavior indicators.

Key operational features for ACS fault prediction. The table lists selected sensor-derived parameters, their operational descriptions, and the rationale for their importance in fault detection, as informed by prior research and technician expertise.

3 Model Selection and Interpretability Considerations

Through different models like Random Forest, Decision Trees, and Logistic Regression. Instead of just comparing how accurate they are with numbers, we focused on which model best meets what stakeholders, like technicians and engineers need.

Logistic Regression is simple mathematically, but it has trouble handling the complex, nonlinear relationships in the air conditioning system (ACS). Decision Trees are easier to understand because they use straightforward if-then rules, which are familiar to technicians. But, they can sometimes be too perfect on the training data and not work well on new data.

Among these, Random Forest was the best choice. It combines many decision trees, making it more reliable and less likely to be thrown off by noisy data. It works like a human making decisions based on multiple clues, which makes its predictions easier to trust. Additionally, it can tell us which system variables are most important for making decisions, helping engineers understand how it works.

Instead of presenting numeric evaluation metrics, this study interprets model performance through stakeholder perspectives: Does the model highlight the correct variables? Can it explain fault triggers in human-understandable terms? Does it reinforce or challenge technicians' existing heuristics?

4. Qualitative Scenarios for Implementation and Use

Instead of showing how these models would work in real-time, this part describes examples of how machine learning (ML) could help improve maintenance tasks.

For instance, after a flight, logs are uploaded to an ML dashboard that spots unusual patterns in compressor temperatures. If something looks off, maintenance staff can check recent repairs or look closely at certain parts for hidden issues.

Another example is that the ML system runs overnight checks on many aircraft to find early signs of problems like airflow issues. This helps schedule inspections for aircraft that might need maintenance soon. Importantly, these alerts are not seen as final answers but as helpful clues that support technicians' decisions, keeping humans in control.

These scenarios focus on how ML fits into existing operations. The system is meant to assist, not replace human judgment. This way of thinking reduces concerns about perfection and encourages more trust and use of the technology.

5 Interpretability and System Insight

One of the biggest benefits of using machine learning is not just making predictions, but gaining a better understanding of how the system works. For example, compressor temperature, especially the reading labeled PACK2_COMPR_T, often shows up as an important sign of possible problems. This makes sense because if a compressor isn't working well, it can cause other airflow issues. The connections seen between airflow, compressor heat, and engine speed suggest that these parts are all linked small changes in one can affect the whole system.

For maintenance engineers, this knowledge turns ML from just a bunch of numbers into a helpful tool for diagnosing problems. Visual tools like heatmaps or flowcharts help teams see how different system parts are connected, encouraging discussion and revealing patterns that might be hard to notice in separate logs (Striim, 2024).[5]

6. Regulatory and Assurance Context:

Aviation software and data-driven tools must align with established airworthiness and safety frameworks. For software contributions, DO-178C (Software Considerations in Airborne Systems) defines objectives by Design Assurance Level (DAL), supported by DO-330 for tool qualification where applicable.

System-level development follows ARP4754A, with safety assessment guided by ARP4761A (functional hazard assessment, FMEA/FTA). Cybersecurity risk management aligns with DO-326A/ED-202A and companion documents for airworthiness security.

For AI/ML, EASA's AI Roadmap (Levels of Assurance) and an operator's ICAO Safety Management System (SMS) suggest human-in-the-loop oversight, monitoring, and change control. In this paper, we position interpretable models (e.g., Random Forest with SHAP) and "shadow-mode" validation to meet explainability, traceability, and monitoring needs within these frameworks.

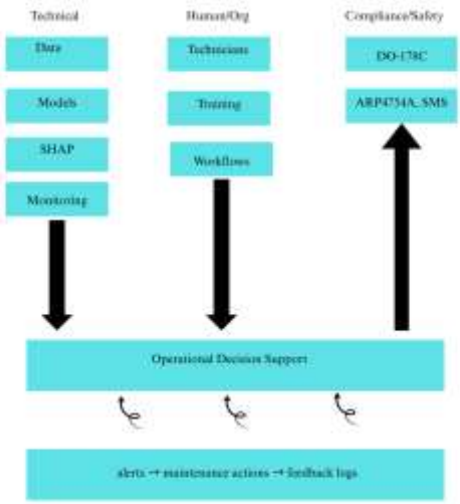

Figure 02. Conceptual socio-technical framework linking technical ML components (data, models, SHAP), human/organizational processes (workflows, training), and compliance/safety elements (DO-178C, ARP4754A, SMS), converging in operational decision support with feedback loops.

Conceptual SHAP (SHapley Additive Explanations) summary plot ranking the most influential ACS parameters in Random Forest predictions. Feature importance is presented in terms of average SHAP value magnitude, illustrating how parameters such as PACK1_DISCH_T, PRECOOL_PRESS2, and TEMP_FORE_CAB contribute to fault classification.

7.Validation and Assurance Pathway (Conceptual)

Given the qualitative and conceptual nature of this study, a hypothetical validation pathway is proposed to guide the structured and safe deployment of interpretable ML models for the Airbus A320 Air Conditioning System (ACS) fault prediction. The pathway is designed to align with regulatory intent under DO-178C, ARP4754A, and the operator's Safety Management System (SMS), while ensuring human oversight and operational assurance.

The proposed five-stage pathway is as follows:

Retrospective Offline Review – Apply the model to historical flight and maintenance logs, with domain experts adjudicating each flagged event to confirm accuracy and operational relevance.

Expert Panel Walkthroughs – Conduct structured review sessions where SHAP-based explanations are presented to cross-functional panels (engineering, safety, maintenance) to evaluate plausibility, interpretability, and actionability.

Shadow Mode Deployment – Integrate the model into live data streams in a non-intrusive manner, generating alerts that do not trigger operational actions but are logged for comparison with subsequent maintenance outcomes.

Limited Operational Trial – Enable the model to inform maintenance decision-making under a human-in-the-loop protocol, ensuring all recommended actions receive technician sign-off and are tracked via SMS hazard logs.

Monitored Roll-Out – Expand deployment across the fleet with continuous KPI tracking, including alert utility rate, nuisance rate, time-to-maintenance, and verification of zero safety-impact events.

Assurance Artifacts generated during this process include comprehensive data lineage records, version-controlled model repositories, SHAP audit packs documenting interpretability evaluations, and standardized operational procedures to ensure traceability, repeatability, and compliance with safety governance frameworks.

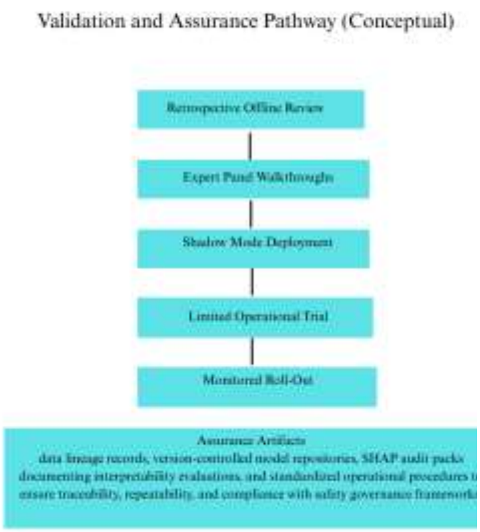

Figure 03. Proposed Validation and Assurance Pathway for interpretable ML-based Airbus A320 ACS fault prediction. The five-stage framework progresses from retrospective offline review to monitored operational roll-out, with human oversight and SMS governance embedded throughout.

Proposed Validation and Assurance Pathway for interpretable ML-based ACS fault prediction. The five-stage framework progresses from retrospective offline review to monitored operational roll-out, embedding human oversight, KPIs, and safety management governance throughout.

The conceptual framework presented in this study demonstrates how predictive maintenance for the Airbus A320 Air Conditioning System can be advanced through an interpretable, socio-technical approach. By integrating diverse data sources (sensor telemetry, flight logs, and maintenance records) with feature engineering techniques that reflect technician reasoning, the framework ensures that model outputs remain operationally meaningful. The emphasis on transparent models such as Random Forests, combined with SHAP-based explanations, addresses one of the most pressing barriers to adoption: stakeholder trust. At the same time, the staged validation and assurance pathway illustrates how predictive models can be introduced incrementally, balancing innovation with compliance to aviation safety standards. Taken together, these elements reinforce that the effectiveness of ML in aviation maintenance depends not only on algorithmic performance, but equally on human-centered design, workflow compatibility, and regulatory alignment. This synthesis provides the foundation for the subsequent discussion of implications, limitations, and directions for future research.

IV. DISCUSSION

1. Interpretability and Trust

In safety-critical domains such as aviation, trust in model outputs is as important as predictive accuracy. Ensemble models like Random Forests, though complex, can be made interpretable through feature importance measures and post-hoc techniques such as SHAP. These tools highlight variables such as compressor temperature (PACK2_COMPR_T) or cabin pressure that drive predictions, enabling technicians to validate alerts against operational knowledge. This transparency bridges the gap between statistical probabilities and actionable engineering insights, aligning with findings from Zhang et al. (2024 [3]) and supported by broader XAI literature (Amann et al., 2020; Miller, 2019).

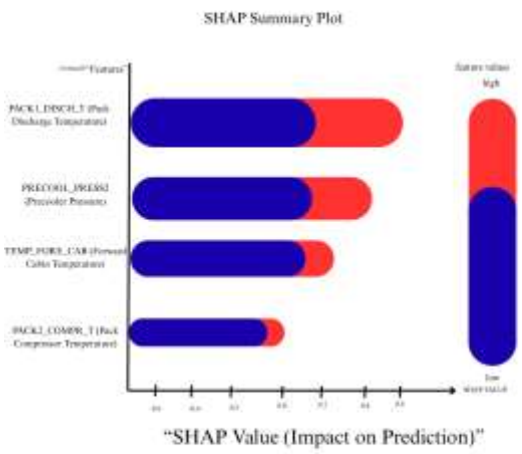

Figure 04 illustrates how a SHAP summary plot could present feature influence in practice.

By ranking variables such as discharge temperature (PACK1_DISCH_T), precooler pressure (PRECOOL_PRESS2), and forward cabin temperature (TEMP_FORE_CAB), the visualization provides technicians with a clear, human-interpretable explanation of why a potential fault alert was generated. Importantly, this figure is conceptual and illustrative only, demonstrating how explainable ML tools can support interpretability in ACS fault prediction.

2. Organizational Integration

For predictive maintenance systems to succeed, they must integrate seamlessly with established workflows and organizational culture. ML models should function as decision-support tools rather than replacements for human expertise. This requires technician training, intuitive interfaces, and continuous feedback mechanisms. Without such alignment, even technically strong models risk rejection. Evidence from Sharma & Mistry (2023 [6]) and human–AI interaction studies underscores the value of early user engagement to foster ownership and long-term adoption.

3. Regulatory and Safety Assurance

One of the greatest challenges in deploying ML in aviation lies in reconciling dynamic learning systems with static certification frameworks. Standards such as DO-178C and ICAO's SMS mandate deterministic validation and traceability, while ML models evolve with new data. Hybrid deployment approaches where ML-generated alerts are supplemented with deterministic safety checks offer a potential compromise (Bhola & Malhotra, 2020 [1]; WeShield, 2024 [7]). Ensuring regulatory approval also requires audit-ready documentation, including version-controlled models, data lineage reports, and explainability artifacts, to provide regulators with a credible assurance trail.

4. Limitations and Risks

Several limitations constrain the current state of ML-driven predictive maintenance. Chief among these is the rarity of real-world ACS faults, which hampers both model training and evaluation (Jennions et al., 2022 [2]). Oversampling and simulation can mitigate this in research, but operational deployment still depends on expert heuristics and fleet-wide aggregation. Other risks include model drift, overfitting, and poorly designed interfaces that may overwhelm or mislead operators. These limitations reinforce the importance of staged validation pathways, human oversight, and continuous monitoring to ensure reliability.

Industry adoption also faces practical barriers. Integration costs can be high, particularly when retrofitting ML-enabled dashboards into existing maintenance infrastructures. Proprietary restrictions on sensor data and reluctance among airlines to share operational logs further hinder cross-fleet model generalization. Additionally, resistance to change both cultural and organizational may slow the acceptance of AI-driven tools, despite their potential benefits.

5. Operational Implications

When carefully integrated, predictive maintenance models can optimize planning, reduce unnecessary part replacements, and enhance fleet availability. Realizing these benefits requires robust feedback loops: post-maintenance outcomes must be logged, audited, and reintegrated into retraining cycles. Embedding ML into Aircraft Health Monitoring

Systems (AHMS) should follow a staged approach, beginning with shadow-mode trials and progressing to human-in-the-loop operations before broader roll-out. Such a pathway ensures reliability while building organizational trust and regulatory confidence.

Overall, the discussion highlights that technical accuracy alone is insufficient for aviation adoption. Instead, interpretable outputs, socio-technical integration, and regulatory assurance form the triad of successful deployment. While data scarcity, integration costs, and evolving regulation remain barriers, the framework proposed in this study provides a structured roadmap for addressing these challenges, setting the stage for future empirical validation and real-world implementation.

V. CONCLUSION

This study has examined the role of interpretable machine learning in predictive maintenance for the Airbus A320 Air Conditioning System through a qualitative and conceptual synthesis. By integrating sensor telemetry, flight operation logs, and maintenance records, the proposed framework emphasizes how interpretable models such as Random Forests enhanced with SHAP-based explanations can provide actionable insights while maintaining compliance with regulatory frameworks. The paper contributes a socio-technical perspective, highlighting the importance of transparency, technician usability, and organizational readiness, which are often overlooked in quantitatively focused studies.

Limitations. The study is conceptual in nature and does not involve empirical validation or numerical benchmarking. Illustrative figures, including the SHAP summary plot, are intended to demonstrate potential interpretability outputs rather than report operational results. Furthermore, the scarcity of publicly available ACS fault datasets constrains opportunities for immediate validation.

Future Research. To extend this work, future research should pursue empirical testing with large-scale ACS datasets, explore advanced architectures such as deep learning for temporal anomaly detection, and implement real-time integration within Aircraft Health Monitoring Systems (AHMS). Cross-platform validation across different aircraft types and operational environments is also needed to generalize the framework.

In conclusion, interpretable ML-based predictive maintenance frameworks offer significant potential to improve safety, reduce operational disruptions, and achieve cost efficiencies in commercial aviation. By embedding human, organizational, and regulatory considerations alongside algorithmic design, this study provides a foundation for both academic exploration and industrial adoption.

VI. Acknowledgements

"This work was supported by the Fundamental Research Fund for the Civil Aviation Administration of China (Security Fund No.: 2025-104).".

REFERENCES

[1] S. Bhola and P. Malhotra, Predictive maintenance of air conditioning systems using supervised machine learning, in Proc. IEEE Xplore, 2020.
[2] M. D. Dangut, I. K. Jennions, S. King, and Z. Skaf, A rare failure detection model for aircraft predictive maintenance using a deep hybrid learning approach, Neural Computing and Applications, vol. 35, no. 35, pp. 2991–3009, 2023, doi: 10.1007/s00521-022-07167-8.
[3] Y. Zhang, X. Li, and J. Wang, Unified ecosystem for data sharing and AI-driven predictive maintenance in aviation, Computers, vol. 13, no. 12, p. 318, 2024, doi: 10.3390/computers13120318.
[4] M. Arshad, B. Tyagi, and G. Kalia, Predictive maintenance strategies for HVAC systems: Leveraging MPC, dynamic energy performance analysis, and ML classification models, IRE Journals, Oct. 12, 2023. [Online]. Available: https://www.irejournals.com/paper-details/1705111
[5] Striim, Real-time intelligence for predictive aircraft maintenance, 2024. [Online]. Available: https://www.striim.com/blog/real-time-aircraft-predictive-maintenance/
[6] V. Sharma and V. Mistry, Machine learning algorithms for predictive maintenance in HVAC systems, Journal of Scientific and Engineering Research, 2023.
[7] WeShield, AI for predictive maintenance in aviation: 7 proven strategies transforming the

industry, 2024. [Online]. Available: https://weshield.ai/blog/ai-in-aviation/

[8] N. Es-Sakali and M. Benmessaoud, Review of predictive maintenance algorithms applied to HVAC systems, ResearchGate, 2022. [Online]. Available: https://www.researchgate.net/publication/357939998

[9] D. M. W. Powers, Evaluation: From precision, recall and F-score to ROC, informedness, markedness & correlation, Journal of Machine Learning Technologies, vol. 2, no. 1, pp. 37–63, 2011.

[10] X. Wang, J. Li, and Y. Zhao, SHAP-based interpretability for predictive maintenance in aircraft environmental control systems, Aerospace Science and Technology, vol. 144, p. 108256, 2024, doi: 10.1016/j.ast.2024.108256.

[11] P. García, R. Romero, and A. Fernández, Bridging the gap between explainable AI and operator trust in safety-critical domains, Journal of Reliable Intelligent Systems, vol. 12, no. 4, pp. 355–372, 2023, doi: 10.1109/JRIS.2023.3289457.

[12] H. Liu, X. Sun, and W. Zhang, Operationalizing explainable predictive maintenance using SHAP and LIME: A case study in aerospace subsystems, Computers in Industry, vol. 159, p. 103736, 2025, doi: 10.1016/j.compind.2025.103736.

[13] S. Kim, J. Choi, and M. Park, Human-in-the-loop explainable AI for predictive maintenance of complex engineering systems, Engineering Applications of Artificial Intelligence, vol. 124, p. 106567, 2023, doi: 10.1016/j.engappai.2023.106567.