FINDING COMMON GROUND FOR AIRBORNE MACHINE LEARNING CERTIFICATION

Emanuele Bezzecchi, Leonardo Helicopters, Varese, Italy
Louis Fabre, Airbus Helicopters, Marignane, France
Darren Cofer, Collins Aerospace, Minneapolis, USA

Abstract

Machine-learning (ML) and, more broadly, artificial intelligence (AI) have become ubiquitous in the consumer market by enabling the development of complex functions inferred from data rather than hard-coded from requirements. Aviation is now adopting ML to address challenges such as perception in degraded-visual environments and optimal flight-path prediction, that conventional engineering cannot readily solve, or that ML can implement more efficiently. Before ML can be fielded in safety-critical aeronautical systems, however, both technical and procedural barriers must be overcome. Regulators in the United States and the European Union are developing complementary, yet presently different, policies, guidance and verification methods to support this transition. This paper compares the emerging certification frameworks of the FAA and EASA, highlights areas of convergence and divergence, and proposes recommendations that may facilitate either formal harmonization or, where necessary, pragmatic coexistence of the two approaches.

1. NOTATION

Acronyms:

ERF	European Rotorcraft Forum
VFS	Vertical Flight Society
AC	Advisory Circular
AMC	Acceptable Means of Compliance
AFE	Authorization for Expenditure
AVSI	Aerospace Vehicle System Institute
DAL	Design Assurance Level
EASA	European Aviation Safety Agency
EU	European Union
EUROCA	European Org. for Civil Aviation
E	Equipment
FAA	Federal Aviation Administration
SAE	Society of Automotive Engineers
TSO	Technical Standard Order
RCL	Recommended Cruise Level
MOC	Means Of compliance
AC	Advisory Circular
AMC	Applicable Mean of Compliance
AL	Assurance Level
IDAL	Item Design Assurance Level
MLC	Machine Learning Constituent
MLDL	Machine Learning Development
IMA	Lifecycle
BASA	Integrated Modular Avionics
ATR	Bilateral Aviation Safety Agreement
	Automatic Target Recognition
RTCA	Radio Technical Commission for Aeronautics

2. INTRODUCTION

Machine-learning technology challenges the foundations of current aeronautical certification frameworks because it departs fundamentally from *requirement to line of code* engineering practice. Contemporary standards such as EUROCAE ED-12C / RTCA DO-178C, EUROCAE ED-79A / SAE ARP4754A, and ED-135 / ARP4761 are grounded in the assumption that system behaviors are explicitly specified, traceable, and verifiable through deterministic testing. ML systems, by contrast, learn behaviors from datasets and produce probabilistic outputs that reflect statistical patterns in the training corpus rather than predefined rules. These outputs are neither **random** nor **unpredictable**; in fact, they are **statistically consistent** with the data distribution on which the model was trained. Given similar inputs, the model will respond consistently, but its logic derives from learnt correlations, not from human-authored requirements.

2.1 Challenges of ML-based systems

This paradigm shift introduces a series of assurance concerns and certification hurdles:

Requirements definition: Behavioural intent for ML-based systems emerges implicitly from training rather than from “shall” statements, complicating the traditional requirements cascade.

Traceability: Numerical parameters—particularly the weights found in deep neural networks—are not human-interpretable. This limits the ability to capture direct trace links between requirements and implementation.

Verification adequacy: Classical structural coverage metrics for software testing lose their discriminatory power because all neurons in a neural network are exercised to some degree on every input. 100% coverage can be obtained with a single test case.

Robustness and generalization: Detecting unintended behavior under edge-case conditions is challenging, as is demonstrating acceptable performance on data outside of the training distribution.

Deployment divergence: Training is performed offline using high-performance tooling and languages that may not be suitable for deployment in safety-critical embedded systems. During operation, model inference is executed on resource-constrained embedded avionics platforms. Differences between environments can introduce subtle discrepancies unless rigorously controlled.

These concerns have been well-documented in both the early AVSI AFE 87 report (Ref. 1) and the following EUROCAE ER-022 /SAE AIR6988 report (Ref. 2), which concluded that traditional development assurance methods are not sufficient on their own to provide an adequate level of confidence for being used in aviation.

2.2 Divergent Regulatory Approaches

Regulators on both sides of the Atlantic also acknowledge these challenges in their own research or proof of concepts but have adopted different strategies to address them.

The European Union is adopting a top-down, standards-first approach.
EASA’s AI Roadmap 1.0 (Ref. 3), released in 2020, and 2.0 (Ref. 4), released in 2023, outline a phased plan that prioritizes the creation of usable guidance (overarching policy, terminology and assurance objectives) before approving the first AI/ML applications.

The United States is tackling the gap with a bottom-up, applications-first strategy.
The FAA began formal engagement on the topic slightly later (in approximately 2022) and emphasizes learning from concrete applications submitted through the Issue Paper process defined by AC 20-166A (Ref. 5). The FAA 2024 Roadmap for AI Safety Assurance (Ref. 6) focuses on defining artefacts and evidence necessary for specific use cases, with broader standardization to follow.

Subsequent sections describe each approach in more detail and assess their implications for industry.

3. EU: TOP DOWN, STANDARDS FIRST

3.1 EASA Policy Context and Guiding Principles

The EU AI Act (Ref. 7) promotes *“…the uptake of human centric and trustworthy artificial intelligence (AI) while ensuring a high level of protection of health, safety …”*. It classifies AI systems by risk and mandates stringent requirements for high-risk domains, including aviation, to ensure a high level of trustworthiness. EASA anticipated in its roadmaps (Ref. 3 and Ref. 4) the **AI trustworthiness** concept as a key enabler to allow acceptance of AI. This concept is based on three technical building blocks (AI Assurance, Human Factors for AI, AI Safety risks mitigation) and is further refined into technical concept papers acting as guidance to support early certification application.

3.2 EASA Guidance Key Technical Considerations

The EASA’s concept papers 1.0 (Ref. 8) and 2.0 (Ref. 9) were established as part of the phase I “*exploration and first guidance development*” of EASA’s **AI Roadmap**. During this phase, several early joint EASA / Industry research projects were launched to elaborate the **concept papers** and the related technical considerations that applicants must address to meet specific AI assurance objectives.

To manage AI applications in aviation, EASA has defined a **classification scheme for AI/ML systems** that groups applications into **three AI Levels** reflecting different tradeoffs of autonomy and human control. Levels 1A and 1B are called *human assistance or augmentation*, levels 2A and 2B are called *human–AI teaming*, and levels 3A and 3B are called *advanced automation*. These levels are useful to scope the guidance: initial policies focus on Levels 1 and 2 for near-term applications with a human in the loop, while Level 3 is slated for future guidance of more autonomous AI once the fundamentals are proven.

In the meantime, EASA enables early AI application using *soft law* tools like **Special Conditions** and **Certification Review Items (CRIs)** based on the concept paper guidance objectives and anticipated means of compliance. This approach has already been applied to initial AI certification applications from several OEM and system/equipment suppliers, with the first formal approvals expected by the end of 2025.

A cornerstone in safety and assurance terminology for AI is the concept of **Learning Assurance**, defined as:

“All those planned and systematic actions used to substantiate, at an adequate level of confidence, that errors in a data-driven learning process have been identified and corrected such that the AI/ML constituent satisfies the applicable requirements at a specified level of performance, and provides sufficient generalisation and robustness capabilities.” (Ref. 9)

In essence, learning assurance extends traditional development assurance to cover the ML-specific aspects of data training and model performance. Together with two other topics, **development and post-operations explainability** and the **data recording capabilities,** they constitute the so-called **AI Assurance** building block.

AI Explainability is defined as “*Capability to provide the human with understandable, reliable, and relevant information with the appropriate level of detail and with appropriate timing on how an AI/ML application produces its results.”* The distinction between development explainability for engineers or auditors and operational explainability for end-users (e.g. pilots, ATM/ANS controllers…) is defined.

In order to support AI/ML explainability, specific considerations are described with regard to **AI Data Recording** for the purposes of monitoring the safe usage of the AI based system, for continuous safety assessment by the applicant, and for accident and incident investigations.

These building blocks and related topics form a new lexicon for systems certification when AI is considered in their design.

The EASA AI Senior Program Manager explains, *“the main difficulty with ML is its disruptive nature, which imposes a change from traditional system, software, and hardware development assurance methods towards an assurance on the learning”* (Ref. 10)

“Consequently, EASA will initially accept only applications where AI/ML constituents do not include IDAL A or B […] items” and limited to those that are developed using **offline supervised learning.**

The result is that due to limited operational experience with ML and lack of accepted assurance techniques, certain safety-critical implementations will not yet be considered by EASA. Even for less critical systems, no credit or assurance level reduction is allowed simply because an ML component is a subset of a larger system. The ML element must be treated with full rigor of its assigned level, rather than assuming a lower criticality within a higher-level system.

These constraints will be revisited as assurance techniques improve, but they reflect the current consensus that ML technology needs more proven reliability before it can shoulder the highest safety burden.

3.3 EASA Augmented assurance process

EASA’s concept papers (Ref. 8, Ref. 9) introduce **Learning assurance processes** and related **Objectives** for ML and outline **anticipated means of compliance (MOC)** to meet new AI assurance objectives. The purpose of these concept papers is to share EASA expectations for approval of ML-based systems as industry consensus practices and standards are being developed. The most prominent concepts are:

Separation of training and inference: the training process takes place offline and is explicitly separated from inference, which is the runtime execution of the trained model on the aircraft or ground system. The training environment (hardware, software libraries, etc.) must be identified and, if certification credit is claimed, qualified accordingly (Objective LM-03). However, the preferred approach is to minimize reliance on this environment. The focus is on the static, configuration-controlled inference model and its integration. Any transformations from training to inference environment, must be documented and demonstrated not to degrade the model’s behavior (Objective LM-06)

Hardware/Software Deployment Constraints: In addition to the separation of training and inference, the hardware and software platform used for inference must remain consistent with the environment validated during certification. Any change to the inference platform or development environment requires an impact analysis to determine which assurance objectives need to be reassessed. Prior certification does not automatically carry over if such changes occur.

Explainability requirements: Any safety-related AI application must provide a level of explainability that allows humans to understand its outputs or recommendations to an appropriate extent. This has two dimensions: developers and certifiers need technical explainability to trace and justify the model’s behavior for debugging and assurance purposes, while end-users and operators require operational explainability to interpret the AI's outputs and support their decisions. For Level 1 and 2 AI applications, where humans remain in or on the loop, operational explainability is especially critical to ensure that end users can trust and effectively use AI assistance. This may include selecting more interpretable architectures, applying explainability techniques (such as heatmaps for vision systems or rule extraction for neural networks), and clearly documenting how outputs relate to inputs in safety analyses. When full insight into a model's internal logic may not be achievable, additional safety measures are expected to mitigate residual risks.

Robustness and performance assurance: It demonstrates that the AI system is robust and consistently meets its performance targets under all foreseeable conditions. ML models can be sensitive to variations or unfamiliar scenarios beyond their training data. Applicants must define suitable performance metrics and validate the model against these metrics on independent datasets that reflect operational conditions (Objectives LM-02 and LM-05). This is linked to the evidence of the model's ability to generalize: Objective LM-04 calls for *quantifiable generalization bounds*. This typically involves a mix of extensive empirical testing on diverse and challenging datasets, along with theoretical analyses like complexity measures or robustness guarantees, to show the model will not fail under new inputs. Robustness also includes the resilience to small input perturbations or noise and behavior in adverse or off-nominal situations.

Traceability and data governance: Because ML development is inherently data-driven, traceability must cover both the data and the learning process setting rigorous data management objectives: datasets for training, validation, and testing must be carefully curated and documented to ensure they suit the intended operational domain, avoiding critical gaps or biases. Data quality requirements, covering completeness, accuracy, labeling correctness and more, must be clearly defined and the data verified accordingly (Objectives DM-07 and DM-08). The entire ML pipeline should be documented to allow tracing how each specific model version was produced, from raw data collection and preparation through training runs to the final deployed model. This traceability enables engineers and auditors to identify root causes of anomalies. It also supports robust configuration management, ensuring that the delivered model corresponds precisely to a known training setup and dataset.

Acceptable ML methods and tools: To manage risk, the EASA concept papers (Ref. 8, Ref. 9) deliberately focus on well-established techniques, primarily off line supervised learning, as these are widely used and better understood in industry. Unsupervised and reinforcement learning are acknowledged but considered more complex and thus not yet fully addressed. If applicants wish to claim assurance that the training process reliably produces a safe model, the training tool environment should be qualified (Objective LM-03). Verification toolchains for AI can include advanced techniques, with EASA actively monitoring developments through initiatives like the joint EUROCAE WG-114 / SAE G-34 committees on AI certification standards, ISO/IEC SC42, and CEN-CENELEC JTC21. Emerging methods such as formal methods for analysis of neural networks, runtime monitors, and automated test data generation, are being explored as acceptable means of compliance. The MLEAP project directly supports this by investigating practical ways to demonstrate, for example, the *quantifiable generalization bounds* required by Objective LM-04, and by assessing verification toolchains for feasibility.

In summary, EASA's current stance is **technologically neutral but pragmatically conservative**. To include all these concepts, EASA's guidance extends the classical V-development model with a mirrored learning V, forming a **W-cycle**. The upper V covers system design and verification; the lower V covers data management, model training and continuous validation. Collectively, these measures aim to deliver safety confidence equivalent to that achieved under conventional processes, for instance EUROCAE ED-79B/SAE ARP4754B and EUROCAE ED-12C / RTCA DO-178C for Software based systems.

4. US: BOTTOM UP, APPLICATIONS FIRST

4.1 FAA Context and Incremental Strategy

The U.S. Federal Aviation Administration (FAA) has adopted a gradual, project-driven approach to AI in aviation.

The FAA began examining ML-based aviation systems around 2021–2022, a few years after EASA's initial steps. In 2022 and 2023, it organized a series of Technical Exchange Meetings with industry to assess the state of the art and understand the needs of early AI applicants. In 2024, insights from these meetings fed directly into shaping the FAA's AI safety roadmap (Ref. 6). Rather than immediately issuing new rules, the strategy is to leverage existing certification processes and learn from real projects in an incremental manner. Applicants are encouraged to introduce AI/ML functions within the current regulatory framework, using existing safety processes except where clearly inadequate. Where novel methods are needed, the FAA will address them through case-specific means of compliance. In practice, this means using the Issue Paper (IP) process (per FAA AC 20-166A) to resolve novel AI certification issues on a case-by-case basis.

4.2 FAA AI Safety Roadmap

The FAA is explicitly pursuing both *the safety of AI and the use of AI for safety* (Ref. 6, page 3). This means ensuring AI systems are safe while also leveraging AI to further improve aviation safety. The document is designed as a *living strategy* that will evolve with technological advances, standards development, and growing operational experience. At its core, the roadmap introduces seven guiding principles that shape how AI certification requirements and methods will develop:

Work within the aviation ecosystem: Integrate AI into the existing aviation safety framework. New approaches are to be considered only when existing methods are clearly insufficient, reflecting confidence in established systems engineering, risk management, and testing practices.

Focus on safety assurance and enhancements: AI should be adopted only if it can meet strong safety assurance criteria and preferably enhance safety or efficiency beyond current capabilities. This principle underscores that verifying AI is safe is just as important as realizing any safety benefits it may offer.

Avoid personification of AI: The FAA cautions against anthropomorphizing AI. Accountability must rest with system designers and developers. Certification will require showing that AI decisions are traceable to requirements and can be understood or overseen by humans.

Differentiate "learned" vs. "learning" AI: The FAA makes a clear distinction between static, trained models ("learned" AI) and adaptive systems that continue learning in operation. Initial approvals will focus on static ML inference models that are not updated once they are deployed.

Take an incremental approach: The FAA advocates introducing AI step by step, starting with narrow, lower-criticality applications (e.g. DAL-D advisory functions) and expanding as confidence grows. This allows the industry and regulator to learn and refine assurance methods based on actual data.

Leverage the safety continuum: Lower-risk environments like experimental aircraft, UAVs, or cargo operations serve as ideal testbeds to gain early experience with AI.

Rely on industry consensus standards: Recognizing the rapid pace of AI advances, the FAA promotes active participation in standards bodies (RTCA, EUROCAE, SAE, ISO, etc.). Consensus standards can adapt more rapidly to technology changes and help harmonize global regulatory requirements, providing applicants with predictable means of compliance.

4.3 Issue Paper Process: Early ML Application

To explore how AI/ML systems can be certified under current regulations, the FAA has encouraged industry to propose pilot projects using case-specific Issue Papers, in line with FAA AC 20-166A. A notable example is Collins Aerospace's *Recommended Cruise Level* (RCL) function, a simple pilot advisory tool that suggests optimal altitudes, implemented via a compact neural network instead of a computationally intensive algorithm. This function was intentionally selected as a low-criticality trial, classified DAL D, since it only provides recommendations with the pilot always in the loop. This allowed both Collins and the FAA to focus on the new types of evidence needed to certify an ML-based system without exposing operations to high safety risk.

To start the Issue Paper process, Collins submitted a white paper summarizing the ML component, including:

Neural network design and intended use: an overview of the architecture and its Operational Design Domain (ODD), detailing the conditions in which it is expected to operate.

Training data and process: descriptions of the datasets, how they were collected and curated, and the training procedure, including tools and parameter choices.

Verification and performance evidence: test results showing that the trained model meets its requirements in terms of accuracy, reliability, and compliance with the advisory function's criteria.

This comprehensive documentation enabled the FAA to evaluate the ML system within the existing certification framework. Through the Issue Paper process, the FAA and Collins agreed on a Method of Compliance demonstrating how these new forms of evidence could satisfy regulatory requirements for airworthiness. The agreement was finally documented as an approved Issue Paper (special project number SP06883WI-Q) documenting this means of compliance for the neural network implementation of RCL and applicable to several TSOs.

Notably, the neural network did not receive any relief in assurance just because it was a software subcomponent. This reflects the FAA's cautious stance that, until broader experience is gained, AI must meet the same safety benchmarks as equivalent traditional systems. The RCL project thus served as a valuable pathfinder, demonstrating that novel ML algorithms can be certified under existing processes, provided applicants supply robust supporting data and analyses.

In summary, the FAA's roadmap does not yet prescribe detailed certification rules but instead lays out a **pragmatic, evidence-driven approach**, focused on developing acceptable MOC. This involves continued research into AI verification methods, pilot projects with industry to put certification processes into practice, and the progressive refinement of guidance material. Over time, this cautious, stepwise approach is expected to align with EASA's framework, jointly shaping robust, harmonized global standards for certifying AI in aviation.

5. HARMONIZATION SCOPE

5.1 ED-324 / ARP6983, Advancing Global Harmonization for AI/ML Certification in Aviation

Even though the EU and the U.S. are pursuing different regulatory paths for ML technologies, both sides share a clear view of the technical and certification gaps that must be bridged. EASA and the FAA have endorsed the concerns raised in **EUROCAE WG-114/ SAE G-34's Statement of Concerns** and the **AVSI AFE 87** study. While their current strategies diverge, both authorities expressly acknowledge the need for a new guidance to evaluate and approve ML-based airborne systems.

The current reciprocal approvals and mutual recognition between leading regulators, mandatory requirements for aviation operating in a single worldwide airspace, are traditionally based on harmonised jointly authored assurance standards that each authority later embeds in its own certification framework. EUROCAE ED-12C / RTCA DO-178C is the classic example. Crafted by an international committee, the standard is accepted by both the FAA and EASA through their respective ACs and AMCs as the definitive means of compliance for the software aspects of airborne systems.

To establish such an assurance standard for AI/ML, the EUROCAE and SAE standardisation bodies tasked a joint working group WG-114 / G-34 to define a Machine Learning Development Lifecycle **(MLDL)** for aviation applications. The proposed standard, EUROCAE ED-324 / SAE ARP6983, is currently a work in progress and subject to ongoing development, review, and revision. The concepts, data, and conclusions presented herein are preliminary and may evolve as technical discussions continue. Due to the complex and technical nature of standards development, the only official and enforceable version of the standard will be the final, formally published document. This paper is provided solely for informational purposes and should not be relied upon as a definitive statement of the final standard or its requirements.

5.2 From Statement of Concerns to Harmonized Guidelines

Prior to launching the development of ED-324 / ARP6983, EUROCAE and SAE documented in the *EUROCAE ER-022 / SAE AIR6988 "Artificial Intelligence in Aeronautical Systems: Statement of Concerns"* the concerns and challenges that the joint working group WG-114 / G-34 will have to address.

Similarly, to the EASA and FAA roadmaps, the joint authored statements of concerns highlights several critical challenges that existing standards and processes are not fully equipped to handle:

Inadequacy of current development assurance approaches: Traditional standards at system and item level (ED-79B/ARP4754B, ED-12C/DO-178C, DO-278A/ED-109A, ED-80/DO-254, ED-153, ED-135/ARP4761A) are not designed to address data-driven specifications, probabilistic design techniques, or the uncertainties they introduce.

Data-driven system/item development: ML models defined by data challenge conventional practices like requirements capture, traceability, and explicit design representation.

Sufficiency of training data: The quality of an ML model is heavily dependent on the data quality (completeness, correctness, representativeness) of the training dataset; gaps can lead to unintended behaviors.

Minimizing unintended behaviors: Due to the inherent nature of data-driven development, unexpected behaviors are harder to predict and detect using conventional validation and verification methods.

Explainability: ML systems may exhibit behaviors that are difficult for designers and end-users to interpret or justify, complicating assurance and acceptance.

Integration with conventional systems: The probabilistic measures typically used to characterize ML-based systems introduce uncertainty that must be accounted for in interactions with traditional deterministic systems, including in areas like safety, operational monitoring, and behavior logging.

To tackle the unique challenges posed by ML in aeronautical systems, the future standard considered insights from other safety-critical domains, including the DEEL white paper for aerospace/rail/automotive (Ref. 11), AMLAS for autonomous systems (Ref. 12), and UL 4600 for automotive autonomy (Ref. 13), to introduce the following foundational concepts:

The "W-Shaped" development path: Similar to the "W-shaped" model in EASA's guidance, this approach maps ML development and integration onto the established aeronautical waterfall lifecycle. It extends the classic "V-shape" by introducing a first "V" dedicated to the ML-specific activities such as model definition, training, validation, and verification and a second connected "V" that covers the implementation and integration of the ML component into conventional systems or items. This ensures ML development is systematically linked back to overall system requirements and verification.

Machine Learning Constituent (MLC): Also aligned with EASA terminology, this concept is used within the SAE ARP4754B / EUROCAE ED-79B framework to identify the combination of ML Models and their associated input/output data processing that perform a specific function. It treats this grouping as a distinct engineering entity for verification and integration within the broader system architecture.

Machine Learning Development Lifecycle (MLDL): The MLDL encompasses the processes and activities specific to the ML Constituent.

Learning Assurance: This concept, also featured in EASA's guidance, builds on the W-shape, MLC, and MLDL to establish confidence that errors in the ML development process have been identified and corrected. It ensures that the MLC meets its applicable requirements at the required performance level. Learning Assurance is essentially the tailored equivalent of classical development assurance, adapted to the data-driven nature of ML.

5.3 From Harmonization Concepts to Harmonized Activities

The W-shape provides an end-to-end mapping from the system level (as defined by EUROCAE ED-79B / SAE ARP4754B in the airborne domain), down through the implementation levels (EUROCAE ED-12C / RTCA DO-178C for software, EUROCAE ED-80 / RTCA DO-254 for hardware), and back to system-level verification. This future standard introduces an intermediate engineering layer dedicated specifically to ML, much like EUROCAE ED-124 / RTCA DO-297 does for IMA systems. It is elaborating process recommendations covering dataset quality and management, the ML model training process, model verification requirements, and implementation and integration activities.

Along the "W-Shaped" development path, the following key processes are defined:

System and safety considerations: This process bridges the ML development with conventional system-level safety assessment activities, which largely remain unchanged. However, due to the unique characteristics of ML, specific interpretations and additional considerations must be introduced to ensure system safety objectives continue to be met.

Machine Learning Development Lifecycle (MLDL): The MLDL comprises several core processes dedicated to the ML Constituent, in particular:

MLC Requirements & Architecture: Defines the functional requirements allocated to the ML Constituent and establishes its overall logical architecture.

ML Data Management: Addresses all aspects related to data, including quality criteria, collection methods, preparation steps, and the organization of data into training, validation, and test datasets.

ML Model Design: Covers the actual development of the ML model, including building, training, and optimization of the ML Models.

ML Validation: Ensures that the requirements and logical *architecture* defined for the MLC are correct, complete, and suitable for their intended purpose.

ML Verification: Confirms that the outputs of the MLDL such as datasets, data processing descriptions, and the trained ML model, meet their specified requirements and detects any errors introduced during the process.

ML Constituent implementation and verification: This stage bridges with the established item-level development standards (EUROCAE ED-12C / RTCA DO-178C for software, EUROCAE ED-80 / RTCA DO-254 for hardware), which themselves remain unchanged. Before transitioning to these item-level processes, it must be ensured that the performance characteristics of the ML Model, as designed and trained during the MLDL, are preserved when mapped onto the actual system architecture or software architecture. **This step is crucial because the fidelity of the implemented ML Model relative to the trained model cannot be verified at the level of individual items**; it must be assessed once integrated into the overall system or software context.

The first official release of this standard, now scheduled for mid-2026, is limited to covering IDAL C and D applications. There is no agreement yet on initial guidance for applications at IDAL A or B. **This is the outcome of a consensus-driven international joint working group** that includes the participation of regulatory authorities.

However, both EASA and the FAA have expressed specific concerns regarding this document and its potential adoption:

EASA is mainly concerned about the timeline for EUROCAE ED-324 and whether it will be available in time to support the rulemaking activities under EASA RMT.0742, as required by the EU AI Act. EASA also emphasized the following key points that have to be respected for its recognition:

- The dual applicability for product certification and ATM/ANS is maintained,
- The end-to-end technical scope for ML constituents and models is maintained,
- The content shall reflect a wide industry range of anticipated practices, based on a multiple use cases and experience gathered so far in research, innovation and on-going certification projects.

FAA is concerned about the standard's emphasis on development processes rather than on the evaluation of the deployed artifact itself. In a series of Technical Exchange meetings with industry, FAA has described a set of engineering characteristics that define the ability of a trained ML model to reliably meet its operational requirements:

1. **Precision of the model**: a measurement of how closely the output of the model agrees with the expected output specified in the modeling data
2. **Behavior of the model**: used interchangeably with the functionality (output) of a model. It includes both the intended behavior (the output of a model for inputs within its operational design domain) and unintended behavior (the output of a model when the input is outside the operational design domain).
3. **Expected error**: statistical measurement of the error when the model is used for prediction
4. **Level of Confidence**: this is a quantitative statistical assessment of how good the model output calculation is.
5. **Sensitivity**: the change in the output with respect to an input perturbation.
6. **Data Completeness**: the domain of valid values for the input must be covered with modeling data (to define the intended behavior).

FAA's expectation is that applicants will need to define requirements based on these characteristics for each ML-based system and then provide specific evidence showing how the deployed system meets these requirements.

6. POSSIBLE PATHS FORWARD

It is critical for the aviation industry that regulatory authorities agree on a common process for approval of ML-based systems for aircraft. Lack of agreement would result in duplicated certification efforts and dramatically increased costs that would likely preclude the use of ML technologies and realizing any of the associated benefits.

However, both FAA and EASA have recently displayed their willingness to progress together in the development of regulations for new technologies, for instance AI, stating *"... Other innovations, for example in artificial intelligence, are emerging rapidly, and ... It is more important than ever that international aviation regulators work together to accompany the changes and ensure safety needs are always met*." (Ref.14)

Moving forward, both industry stakeholders and regulators should pursue complementary strategies to bridge the gap between the divergent U.S. and EU approaches. Key recommendations include:

Accelerate common, harmonized standards: Support the finalization and adoption of joint AI/ML certification standards like SAE ARP6983 / EUROCAE ED-324, which are being developed by international committees with FAA and EASA participation. Once published, these standards can serve as de facto harmonized means of compliance, analogous to how DO-178C/ED-12C functions for software today. Regulators should be ready to endorse such consensus standards in their guidance (e.g. via an FAA AC and EASA AMC) so that applicants can follow one globally accepted process for ML assurance.

Leverage bilateral airworthiness agreements: The existing US–EU Bilateral Aviation Safety Agreement (BASA) should be used to recognize and accept each other's certification findings for AI systems whenever possible. In a recent joint statement, FAA and EASA pledged to *"fully leverage the US-EU Safety Agreement"* to reduce duplicated efforts and *"deepen proactive collaboration on certification activities for new and innovative technologies"* (Ref. 14). Building on this spirit, the authorities could establish agreed-upon special conditions or technical issue papers for AI that are mutually acceptable.

Joint certification pilot projects: Initiating projects with combined FAA/EASA engagement can provide practical experience and foster trust. Early collaborative reviews would help identify differences in interpretation, validate new verification tools, and ultimately converge on mutually agreeable acceptance criteria. Additionally, research initiatives like EASA's MLEAP or FAA's Tech Center studies should share findings openly, so lessons learned on one side inform the other.

Address near-term gaps pragmatically: Given that full harmonization will take time, regulators should provide clear interim guidance to industry on how to navigate the differences in the meantime. This may include publishing more detailed policy Issue Papers/Certification Review Items (CRIs) for common AI use-cases, so that expectations are transparent. Industry, for its part, should be prepared for extra diligence, developing ML components in a modular way that can satisfy the most stringent requirements from either side.

Without alignments in the short term, applicants may need to duplicate certification efforts, driving up cost and time to market. Such duplication risks delaying the introduction of safety-beneficial AI technologies or making it economically infeasible to field them. A cooperative stance will mitigate unnecessary friction until formal rules catch up. By pursuing these steps, the aviation community can gradually build a common safety assurance ecosystem for AI. The goal is to reach a point where an AI-based system can be approved once and accepted globally, just as traditional avionics are today. This will require patience and continued dialogue, but the payoff is a certification approach that upholds safety while enabling innovation across all markets.

7. CONCLUSION

AI/ML technologies offer significant opportunities for aviation, but these benefits depend on the ability to certify and trust such systems in safety-critical roles. In this paper, we examined the two main regulatory approaches: EASA's structured, roadmap-driven framework and the FAA's more incremental, project-based strategy. We also explored how the EUROCAE WG-114 / SAE G-34 committee is developing a new standard through a joint international effort, aiming to establish common practices and build a bridge between these different paths, and laying the foundation for a harmonized global framework.

It is the authors' opinion that, establishing a common ground for certification has a significant impact, contributing to greater efficiency and coherence across the helicopter industry, maintaining aviation's high safety standards while facilitating innovations that bring real value to society and helicopter operators with functionalities like:

- Locating casualties in difficult conditions and support search-and-rescue or disaster relief
- Improving safety and situational awareness to detect unforeseen obstacles through specific sensors, automatically recalculating safe flight paths
- Increasing flight deck automation to reduce crew workload in demanding mission or flight phases
- Improving early detection of systems and mechanical parts failures
- Improving maintenance and Helicopters availability with dedicated on Condition Based Maintenance

This means economic gains for operators, better service reliability, and improved responsiveness in emergency missions.

In short, harmonizing airborne AI aspects of certification is not bureaucracy for its own sake, it is the key to progress. A unified assurance framework keeps safety paramount while allowing aviation to leverage cutting-edge technology. **This future is within reach if regulators and industry continue to work together with a shared commitment to safety and innovation.**

Authors contact:

Emmanuele Bezzechi
emanuele.bezzechi@leonardo.com

Louis Fabre louis.fabre@airbus.com

Darren Cofer darren.cofer@collins.com

8. REFERENCES

1. AFE 87 Project Members, Final Report AFE 87 – Machine Learning, Aerospace Vehicle Systems Institute, Tech. Rep. 87-REP-01, May 7, 2020.
2. EUROCAE/SAE, Artificial Intelligence in Aeronautical Systems: Statement of Concerns ,EUROCAE ER-022 / SAE AIR6988, Apr. 2021.
3. European Union Aviation Safety Agency, Artificial Intelligence Roadmap 1.0: A human-centric approach to AI in aviation, Cologne, Germany, 2020.
4. European Union Aviation Safety Agency, Artificial Intelligence Roadmap 2.0: A human-centric approach to AI in aviation, Cologne, Germany, 2023.
5. Federal Aviation Administration, Advisory Circular 20-166A: Issue Paper Process, Nov. 6, 2014.

6. Federal Aviation Administration, Roadmap for Artificial Intelligence Safety Assurance, Ver. I, https://www.faa.gov/media/82891
7. Regulation (EU) 2024/1689 of the European Parliament and of the Council, laying down harmonised rules on artificial intelligence. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32024R1689&qid=1753171558243
8. European Union Aviation Safety Agency, Concept Paper: First Usable Guidance for Level 1 Machine Learning Applications, 2021 https://www.easa.europa.eu/en/newsroom-and-events/news/easa-releases-its-concept-paper-first-usable-guidance-level-1-machine-0
9. European Union Aviation Safety Agency, Artificial Intelligence Concept Paper Issue 2: Guidance for Level 1 & 2 machine-learning applications, Mar. 2024 https://www.easa.europa.eu/en/newsroom-and-events/news/easa-publishes-artificial-intelligence-concept-paper-issue-2-guidance
10. Adapting regulation to AI – Leading the way with the AI roadmap of the EASA https://www.flyinstinct.com/adapting-regulation-to-ai
11. Dependable and Explainable Learning (DEEL) Certification Workgroup, Machine Learning in Certified Systems, White Paper, IRT Saint Exupéry, Toulouse, France, 2021.
12. Hawkins, R., Paterson, C., Picardi, C., Jia, Y., Calinescu, R., and Kelly, T. P., Guidance on the Assurance of Machine Learning in Autonomous Systems (AMLAS), Assuring Autonomy International Programme, University of York, UK, 2021.
13. Underwriters Laboratories, Standard for Safety for the Evaluation of Autonomous Products, UL 4600 (1st Edition), Apr. 2020.
14. “FAA and EASA Pledge Strong Cooperation to Address Aviation Challenges of the Next Decade,” https://www.faa.gov/newsroom/faa-and-easa-pledge-strong-cooperation-address-aviation-challenges-next-decade#:~:text=,aviation%20fuel%20development%20and%20deployment