More Regulation or Better Guidance? User Perspectives on Accountability in Aviation AI-based Systems

Abstract: This paper explores whether aviation professionals' demand for "more regulation" on responsibility and accountability in AI-based aviation systems should be understood as a demand for additional hard law or more guidance material. Based on EASA's Ethics for AI in Aviation Survey, the LOKI ELSA assessment with air traffic controllers interacting with an AI-enabled controller and EASA's Concept Paper Issue 03, this paper connects macro-level regulatory expectations with micro-level user feedback. The analysis shows that accountability concerns often arise, not necessarily only from legislative gaps, but from gaps in role clarity, explainability, task allocation, meaningful oversight and operational implementation and not necessarily from legislative gaps. By mapping the LOKI findings to EASA's AI guidance provided through the Concept Paper Issue 03, this paper argues that sector-specific guidance can bridge legal requirements and operational practice in a way that is technically robust, proportionate and meaningful for aviation professionals.

Keywords: Aviation AI Governance, Accountability, EU AI Act, Human Oversight, Regulatory Guidance

Author(s): Sofia Kyrampalidou, DLR, sofia.kyrampalidou@dlr.de

Section 1: Introduction

Aviation has traditionally depended on highly trained professionals, formal operational procedures and protocols, as well as a safety-critical decision-making culture in order to be able to manage complex operations.[1] Nowadays it is among the domains that actively explore the benefits of artificial intelligence (AI) across multiple areas, such as air traffic management, aircraft operations, maintenance and airport operations, where AI-enabled systems promise to support decision-making, prediction, optimization as well as partial automation of tasks.[2] Considering the increasing complexity and demands for higher levels of efficiency in aviation, the benefits of this support are surely beneficial.[3] However, aviation is a sector where such solutions cannot be evaluated only in terms of performance or efficiency. New tools must be assessed in terms of safety, human factors, organizational responsibility and the ability of the professionals to understand and control the system they work with.

Accountability is particularly important. When an AI-based system supports or performs parts of an operational task, the allocation of responsibility and accountability among the involved

[1] See Karlene H Roberts, "New Challenges in Organizational Research: High Reliability Organizations" (1989) 3:2 Industrial Crisis Quarterly 111, DOI: <10.1177/108602668900300202>; David A Wiegmann et al, "Safety Culture: An Integrative Review" (2004) 14:2 The International Journal of Aviation Psychology 117, DOI: <10.1207/s15327108ijap1402_1>; Nili Sukenik, "Coordination in the Cockpit: A Game-Theory View of Standard Operation Procedures" (1998) 8:4 The International Journal of Aviation Psychology 405, DOI: <10.1207/s15327108ijap0804_5>.

[2] See European Union Aviation Safety Agency, *Artificial Intelligence Roadmap 2.0: A Human-Centric Approach to AI in Aviation* (Cologne: EASA, 2023) [EASA Roadmap 2.0]; Cristian Lozano Tafur et al, "Applications of Artificial Intelligence in Air Operations: A Systematic Review" (2025) 25 Results in Engineering 103742, DOI: <10.1016/j.rineng.2024.103742>.

[3] See European Aviation High Level Group on Artificial Intelligence, *The FLY AI Report: Demystifying and Accelerating AI in Aviation/ATM* (Brussels: EUROCONTROL, 2020); Lozano Tafur et al, supra note 2.

stakeholders becomes more complex.[4] This becomes even more challenging when the focus is on a Human-AI teaming scenario. When a professional is working real-time with an AI-based assistant that may recommend or even perform some actions under certain conditions, it is no longer sufficient to just ask whether the system performs its tasks only from a technical point of view. It is also important to ask whether the professional understands the system's roles and limitations; whether the organization they act within has defined responsibilities and whether the applicable regulatory framework provides enough guidance in order to make accountability clear in an operational setting.

This paper starts from a pattern identified in recent AI governance discussions. Within the context of the European Union Aviation Safety Agency's (EASA) study on Ethics for AI in Aviation[5] published in 2025, 231 aviation professionals were asked about their ethical considerations and perceptions regarding AI applications in aviation. Their responses indicated a strong demand for regulation across multiple of the hypothetical scenarios introduced by EASA, but with a strong focus on the use case dealing with teaming with an AI-based system. At first, one might think that this indicates the need for more regulation, but this interpretation is a bit narrow. When aviation professionals ask for more "regulation", they do not always ask for new legal obligations, but rather they might be asking for clearer implementation guidelines, such as role allocation, operational concepts and assurance practices that can make accountability understandable and feasible in practice.

The LOKI project[6] provides an insightful micro-level perspective on this question. Within the context of LOKI, 7 air traffic controllers (ATCOs) were asked to interact with a digital, AI-enabled Digital Interactive Reliable Controller (DIRC) and to provide feedback on their experience. The study translated the EU AI Act's high-risk requirements into user-facing questions; not with the main goal to prove compliance, but rather to understand how end users perceive legal requirements in a realistic operational setting and how/if such requirements actually meet the users' expectations when implemented. The results showed conditional acceptance, as the ATCOs were generally very open to the step-by-step introduction of DIRC-like solutions, especially for routine tasks, if some improvements or certifications were made. At the same time, there were also some concerns regarding edge cases and the role and limitations of the system. These findings point towards the claim that accountability concerns do not arise only from gaps in regulation, but also from gaps in operational guidance.

EASA has also recently published their third version of the EASA Concept Paper: Guidance for safety-related Artificial Intelligence applications, as a draft.[7] This issue provides a detailed set of objectives and anticipated means of compliance for safety-related AI applications. It introduces and refines concepts such as trustworthiness analysis, AI assurance, human-AI cooperation and collaboration, responsibility scheme assessment and others. Although the Concept Paper is not

---

[4] See Helen Nissenbaum, "Accountability in a Computerized Society" (1996) 2:1 Science and Engineering Ethics 25, DOI: <10.1007/BF02639315>; Filippo Santoni de Sio & Giulio Mecacci, "Four Responsibility Gaps with Artificial Intelligence: Why They Matter and How to Address Them" (2021) 34:4 Philosophy & Technology 1057, DOI: <10.1007/s13347-021-00450-x>.

[5] European Union Aviation Safety Agency, *Ethics for Artificial Intelligence in Aviation: Aviation Professionals Survey Results 2024/2025* (Cologne: EASA, 2025) [EASA Ethics Report].

[6] Deutsches Zentrum für Luft- und Raumfahrt, "LOKI – Kollaboration von Luftfahrt-OperateurInnen und KI-Systemen" (last visited 23 July 2026), online: <dlr.de> [LOKI website]. The research underlying this paper was conducted as part of the LOKI project (Collaboration of Aviation Operators and AI systems, 2022—2026), funded by the German Aerospace Center (DLR).

[7] European Union Aviation Safety Agency, *EASA Artificial Intelligence Concept Paper — Proposed Issue 3* (Cologne: EASA, 2026) [EASA Concept Paper Issue 03].

a binding regulation, it provides a structured regulatory layer between high-level legal requirements and concrete operational implementation.

The main argument of this paper is that the aviation professionals' demand for accountability should not be directly translated as demand for a new regulation. In some cases, it might be better understood as a demand for more mature regulatory guidance, meaning the translation of the high-level requirements into domain-specific expectations, use cases, responsibility schemes, human-centered design principles and appropriate assurance objectives. The paper contributes to this discussion by mapping the LOKI user feedback against the objectives and concepts found in EASA's updated draft Concept Paper. This mapping should not be perceived as a legal compliance assessment, but rather as a "perception to guidance" tool that shows how user concerns can help regulators understand where guidance – and not necessarily more regulation – is needed.

This paper is structured as follows: Section 2: Regulation, guidance and accountability in aviation presents the current landscape of regulation, guidance and accountability in aviation on which this paper was based; Section 3 discusses the need for regulation as per the findings of the EASA study; Section 4 presents the study that took place within the context of the research project LOKI; Section 5 maps the LOKI findings to the objectives and concepts found in EASA's Concept Paper Issue 03; Section 6 discusses the results and explores this need for more regulation or better guidance and Section 7 concludes the paper.

Section 2: Regulation, guidance and accountability in aviation

Aviation regulation has combined binding rules with detailed guidance, standards and certification processes for a long time.[8] The higher-level rules tend to define obligations and responsibilities linked to safety outcomes while guidance and means of compliance support the involved stakeholders (e.g., airlines, authorities) to better understand how these obligations can be realized in practice.[9] Now, governing AI applications in aviation seems to be following a similar layered approach, but this time additional complexities need to be tackled, as the new AI-based systems challenge what has been established so far regarding predictability, explainability, technical assurance and human-machine/human-AI interaction.[10]

The EU AI Act provides a horizontal legal framework for AI systems, including high-risk AI systems. It establishes requirements that are relevant to risk management, data governance, technical documentation, record keeping/logging, transparency, human oversight, accuracy and robustness.[11] These requirements are also relevant for aviation AI applications, and particularly where they are safety-related or part of a safety-component.[12] However, the AI Act itself is by no means an operational aviation manual. It does not specify things such as: how can an air traffic

[8] See John Woodlock & Håkan Hydén, "(f)Lex Avionica: How Soft Law Serves as an Instrumental Mediator Between Professional Norms and the Hard Law Regulation of European Civil Aviation Maintenance" (2020) 121 Safety Science 54, DOI: <10.1016/j.ssci.2019.08.037>.

[9] Ibid.

[10] See Konstantin Dmitriev, Johann Schumann & Florian Holzapfel, "Toward Certification of Machine-Learning Systems for Low Criticality Airborne Applications" in *40th Digital Avionics Systems Conference, DASC 2021 — Proceedings* (New York: IEEE, 2021), DOI: <10.1109/DASC52595.2021.9594467>.

[11] Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828, [2024] OJ L 2024/1689, arts 8—15 [AI Act].

[12] Ibid, arts 6(1), 8—15, 108, Annex I.

controller know the roles and limitations of an AI-enabled digital assistant or how responsibility should be allocated in a Level 2B human-AI collaboration scenario. These questions need to be viewed through a specific sectoral lens for better translation and implementation.

EASA's AI Concept Paper Issue 03[13], although currently in its draft and open for consultation version, can already be perceived as part of this sectoral lens. It should be read complimentary to the AI Act and other already existing aviation safety rules and it provides a structured framework of objectives and anticipated means of compliance for safety-related AI applications. Its goal is to guide organizations and applicants on how they could introduce such AI technologies in aviation domains covered by EASA's Basic Regulation[14] and to prepare for future rule-making activities. Due to its nature, the Concept Paper appears to be positioned in a middle space between hard law and operational practice.

This layered approach could assist with accountability matters in aviation AI applications, as at the legal level, accountability concerns the allocation of obligations, responsibilities and liabilities among the involved actors.[15] At the organizational level, it concerns the capacity of the organizations to manage risks, their competences, documentation and monitoring or reporting.[16] At the operational level, it concerns whether the human professionals involved understand the AI system in terms of whether they can understand what the system does, what limits it has, when it should be relied upon, when it should be overridden and who is responsible for which task.[17] A regulatory framework that focuses only on one of these levels will most likely be insufficient.

This is particularly relevant for human-AI teaming. In traditional automation scenarios, the main question often focused on whether the human was in the loop or able to supervise the automated functions. In AI-based collaboration or cooperation, however, this mentality might be too simplistic, as the system may recommend, prioritize, allocate or perform part of a task while remaining a "black box".[18] This distinction between cooperation and collaboration is important here because cooperation may involve a looser distribution of tasks, while collaboration may suggest a closer interdependence of tasks. In both cases, a human could have authority on paper, but may also lack sufficient information to exercise meaningful human control, especially under time pressure or overreliance.[19] This creates important accountability problems, not because

---

[13] EASA Concept Paper Issue 03, supra note 7.

[14] Regulation (EU) 2018/1139 of the European Parliament and of the Council of 4 July 2018 on common rules in the field of civil aviation and establishing a European Union Aviation Safety Agency, and amending Regulations (EC) No 2111/2005, (EC) No 1008/2008, (EU) No 996/2010, (EU) No 376/2014 and Directives 2014/30/EU and 2014/53/EU, and repealing Regulations (EC) No 552/2004 and (EC) No 216/2008 and Council Regulation (EEC) No 3922/91, [2018] OJ L 212/1 [EASA Basic Regulation].

[15] See Claudio Novelli, Mariarosaria Taddeo & Luciano Floridi, "Accountability in Artificial Intelligence: What It Is and How It Works" (2024) 39:4 AI & Society 1871, DOI: <10.1007/s00146-023-01635-y>.

[16] See Inioluwa Deborah Raji et al, "Closing the AI Accountability Gap: Defining an End-to-End Framework for Internal Algorithmic Auditing" in *FAT* '20: Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (New York: Association for Computing Machinery, 2020) 33, DOI: <10.1145/3351095.3372873>.

[17] See Luciano Cavalcante Siebert et al, "Meaningful Human Control: Actionable Properties for AI System Development" (2023) 3 AI and Ethics 241, DOI: <10.1007/s43681-022-00167-3>.

[18] See Lisanne Bainbridge, "Ironies of Automation" (1983) 19:6 Automatica 775, DOI: <10.1016/0005-1098(83)90046-8>; Mica R Endsley, "Supporting Human-AI Teams: Transparency, Explainability, and Situation Awareness" (2023) 140 Computers in Human Behavior 107574, DOI: <10.1016/j.chb.2022.107574>; Allyson I Hauptman et al, "Understanding the Influence of AI Autonomy on AI Explainability Levels in Human-AI Teams Using a Mixed Methods Approach" (2024) 26 Cognition, Technology & Work 435, DOI: <10.1007/s10111-024-00765-7>.

[19] See Cavalcante Siebert et al, supra note 17.

there is no legal framework indicating that a human is responsible, but rather because the operational conditions under which the human can exercise meaningful human control are not sufficiently defined in order for them to act responsibly.

The distinction between responsibility and accountability is also important. Responsibility may refer to the tasks/duties that are assigned to an actor in a given operational setting. Accountability refers to the capacity and obligation of the involved actors (high-level, organizational or operational) to answer when necessary for decisions, actions and outcomes – especially after incidents or failures.[20] In AI-based systems, accountability cannot be assigned to an AI system as if it was a legal or a moral actor.[21] Therefore, the relevant question is how accountability is distributed among the actors involved in its design, approval or deployment. In this context, the Concept Paper Issue 03's suggested responsibility scheme is important because it provides an operational basis for accountability: it requires all responsibilities to be allocated for each high-level task/operational scenario, while ensuring that no responsibility is assigned to the AI system itself, no responsibility gaps are left and that responsibility is not shifted to actors that do not have the authority or the information to exercise it. As a result, responsibility allocation becomes a prerequisite for meaningful accountability.

Section 3: EASA AI Ethics Survey – macro-level demand for regulation

The EASA Ethics for AI in Aviation survey provides a broad view of how aviation professionals perceive AI-based systems across different hypothetical use cases.[22] Within the context of this study, aviation professionals were asked to evaluate scenarios that included an AI aspect based on their level of comfort, trust and acceptance.[23] The results showed a cautiously positive, but uncertain attitude towards AI applications in aviation. Although they did not unanimously express negative views, their reported levels of comfort, trust and acceptance were not strong enough to suggest that AI-based systems are already perceived as "approved".

The survey is particularly relevant for this paper, as it shows a strong demand for regulation. The aviation professionals expressed the view that AI-based systems should be regulated across the different use cases that were described (be it safety or non-safety related settings). The demand was stronger in the use cases concerning human-AI teaming and physiological data monitoring. EASA was also indicated as the main authority responsible for oversight, supported by national aviation authorities and other relevant institutions.

At a first reading, this seems to support the claim that aviation AI requires more regulation. However, a closer look shows that the reasons for non-acceptance of these AI-based systems were not only limited to legal uncertainty. Professionals raised some concerns regarding the systems' performance, the definition of the operational domain, operational mitigation measures, mistrust of AI, data protection, accountability, professional threats, human oversight, safety demonstration and the impact on human competence. These concerns are concrete, operational and sociotechnical. They do not all require the same regulatory response.

---

[20] See Mark Bovens, "Analysing and Assessing Accountability: A Conceptual Framework" (2007) 13:4 European Law Journal 447, DOI: <10.1111/j.1468-0386.2007.00378.x>; Santoni de Sio & Mecacci, supra note 4.

[21] See Joanna J Bryson, Mihailis E Diamantis & Thomas D Grant, "Of, For, and By the People: The Legal Lacuna of Synthetic Persons" (2017) 25:3 Artificial Intelligence and Law 273, DOI: <10.1007/s10506-017-9214-9>; Santoni de Sio & Mecacci, supra note 4.

[22] EASA Ethics Report, supra note 5.

[23] Ibid.

The responsibility and accountability findings are also especially important. In the case of human-AI teaming, aviation professionals considered the deploying organization of the human operator to be the most responsible and accountable, while assigning the least responsibility to the human end user. This is important because it counters simplistic views that the human operator that is "in the loop" should remain primarily accountable whenever an AI system is used. If an organization deploys an AI-based system, defines its use, integrates it into daily operations and benefits from its use, then this organization should be casted with responsibility or accountability and not solely the end user.[24]

This view is also consistent with the broader principle of meaningful accountability: responsibility should be assigned to actors that have the authority, the information and the capacity to control an AI-based system.[25] If a human professional is expected to oversee an AI system, but does not understand its operational domain or limits, then assigning responsibility to that professional risks becoming "responsibility dumping".[26] If an organization that deploys such systems controls things like integration, trainings, procedures and monitoring then it should hold the largest share of responsibility ensuring that the systems is safe and transparent.

The report also includes suggestions for future EASA initiatives. Aviation professionals called for standards and guidelines for ethical use of AI, principles on collaboration/cooperation between humans and highly automated systems, guidance on deployment, alignment with existing regulation, stakeholder engagement, expert groups, early involvement of end users, trainings, knowledge exchange and clarity when it comes to certifications. All these proposals suggest that this demand for "regulation" is not solely a demand for legal rules, but rather a demand for a clear regulatory leadership, practical guidance, clear communication and competence building.

This distinction is the core of this paper's thesis. When aviation professionals call for regulation, they may be asking for assurance that AI will not be informally or unsafely introduced. To address this, the regulators should analyze which of the aforementioned concerns require binding rules and which could be addresses through guidance, acceptable means of compliance, certification practices or more explanatory material, such as the Concept Paper.

Section 4: The project LOKI – micro-level user feedback to regulation

Within the context of the LOKI project, we researched and developed fundamentals for introducing AI systems in air traffic and flight control.[27] One of the aspects of the project was the Ethical, Legal and Social Aspects (ELSA) assessment of the AI-enabled Digital Interactive Reliable Controller – DIRC. DIRC is a digital air traffic controller that acts as a highly automated team partner which collaborates with human ATCOs. DIRC can analyze air traffic, take decisions and

[24] See Ida Koivisto, Riikka Koulu & Stefan Larsson, "User Accounts: How Technological Concepts Permeate Public Law Through the EU's AI Act" (2024) 31:3 Maastricht Journal of European and Comparative Law 412, DOI: <10.1177/1023263X241248469>; Madeleine Clare Elish, "Moral Crumple Zones: Cautionary Tales in Human-Robot Interaction" (2019) 5 Engaging Science, Technology, and Society 40, DOI: <10.17351/ests2019.260>.

[25] See Madalina Bryson, Joanna J, Mihailis E Diamantis & Thomas D Grant, "Of, For, and By the People: The Legal Lacuna of Synthetic Persons" (2017) 25:3 Artificial Intelligence and Law 273, DOI: <10.1007/s10506-017-9214-9>.

Busuioc, "Accountable Artificial Intelligence: Holding Algorithms to Account" (2021) 81:5 Public Administration Review 825, DOI: <10.1111/puar.13293>; Cavalcante Siebert et al, supra note 17.

[26] EASA Concept Paper Issue 03, supra note 7.

[27] LOKI website, supra note 6.

implement them without explicit human approval, which classifies it under Level 2B of the EASA AI automation taxonomy.[28]

The ELSA assessment was based on the EU AI Act[29] where the selected high-risk system requirements were translated into user-facing questions regarding risk management, data governance, technical documentation and explainability, record keeping, transparency, human oversight, accuracy and robustness, AI literacy and acceptance. After the final validation testing of the system, the air traffic controllers that were involved in the project activities (ATCOs) (n=7) were asked about their interaction with DIRC.

The main purpose of the ELSA assessment was not to prove compliance with the AI Act as subjective user perceptions cannot constitute legal compliance evidence. Instead, it aimed to understand how compliance-oriented requirements are perceived by the end users and how these requirements actually take form in operational settings. In this manner, we wanted to take a first step towards exploring how ATCOs, also being aviation professionals, perceive legal requirements and what is their feedback on the requirements' adequacy or efficiency. Here are the main findings.

AIA requirement	Brief description	Questions	User response (n=7)	User explanations
Art. 9 – Risk Management System	The system must manage risks across different contexts and avoid creating new risks	Did DIRC ever suggest an action that you felt could increase risk?	Yes: 5 / No: 2 	5 of 7 described an unsafe action of DIRC
		Did you feel DIRC adapted safely to changes in context? (e.g., traffic, fatigue levels, etc.)	Yes: 5 / No: 2 	mainly based on experience in the basis run

Figure 1. LOKI findings regarding EU AI Act, art 9 - Risk Management.

First, users generally thought that DIRC adapted safely to changing contexts, but some of them recalled at least one unsafe suggestion. This shows that the overall trust in the system might be there, but it is hindered by edge cases. In aviation, such cases matter as they could have a great impact on safety outcomes and naturally, trust. The findings also show the value of user feedback as a signal for **risk management**. End users that are actively involved in operations may identify scenarios that technical design and test processes might have fully captured yet.

[28] EASA Roadmap 2.0, supra note 2.
[29] AI Act, supra note 11.

AIA requirement	Brief description	Questions	User response (n=7)	User explanations
Art. 10 – Data & Data Governance **(user-facing: perceived fairness & consistency)**	Data must be relevant, representative, accurate and non-discriminatory. Users won't see the datasets but might experience issues linked to them	Did you notice any biased, strange, or systematically wrong suggestions in specific conditions?	Yes: 0 / No: 5 / Unsure: 2 	possibly even fairer, DIRC a little more cautious

Figure 2. LOKI findings regarding EU AI Act, art 10 - Data Governance.

Since **data governance** is not directly accessible and understood by end users, this requirement was translated into a more user-facing aspect of biased outcomes or strange behavior that could be observed by the end users. ATCOs did not report any obvious bias or systematically strange behavior when using DIRC. However, the study noted that this does not prove data representativeness, non-discrimination or data quality as indicated by the legal requirement. This distinction matters for regulatory guidance, as user perception can contribute to monitoring, but back-end assurance remains needed.

AIA requirement	Brief description	Questions	User response (n=7)	User explanations
Art. 11 – Technical Documentation **(user-facing: explainability)**	Users should be provided with adequate information on the system's logic and limits	Did DIRC provide you with enough information to understand why it made a recommendation/ action?	Yes: 2 / No: 3 / Unsure: 1 / Partly: 1 	overview & history in HI-DIRC known to all, was partially used; feedback loop or knowledge of DIRC's plans was missing for explainability and trust, requirement not fulfilled to being able to understand DIRC's actions without having to delve too deeply into them
		Did you feel that you could trust the explanations provided?	Yes: 2 / No: 1 / Unsure: 3 / Partly: 1 	

Figure 3. LOKI findings regarding EU AI Act, art 11 - Technical Documentation.

The requirement regarding the **technical documentation** that should be kept in file was translated towards the more user-facing notion of explainability. Here ATCOs were split on whether they received sufficient information to understand DIRC's actions and recommendations fully. They also missed a clearer understanding on the system's intended plan or next steps. This is highly relevant for accountability matters. A human controller cannot meaningfully monitor and be responsible for an interaction with an AI-enabled assistant if the assistant's rationale and "what comes next" actions remain unclear at decision time.

AIA requirement	Brief description	Questions	User response (n=7)	User explanations
Art. 12 – Record-keeping / Logging	The system must log events in order to support traceability	Were you aware that DIRC was logging or recording interactions?	Yes: 5 / No: 2 / Unsure: 0 / Partly: 0 	history offered by HI-DIRC, HI-DIRC partially used, not primarily, but for overview if required
		Would being able to review a history of DIRC's actions help you in your work?	Yes: 6 / No: 0 / Unsure: 0 / Partly: 1	

Figure 4. LOKI findings regarding EU AI Act, art 12 - Recording Keeping/Logging.

Logging and history were positively perceived by the end users as they seemed to be aware that DIRC recorded interactions and most of them found the ability to review the history of actions useful. Nonetheless, the analysis also showed that this ability to review historical actions was not fully integrated into the routine workflow, as the ATCOs did not use it consistently. Once again, this shows the difference between having the capability to do something and this specific capability actually being operationally useful. In general, record keeping is not solely a technical compliance function, as it also supports learning, debriefing, traceability and accountability.

AIA requirement	Brief description	Questions	User response (n=7)	User explanations
Art. 13 – Transparency & Information	Users must be informed about the system's capabilities, limitations and overall role	Was it always clear that the recommendations/ actions came from AI (DIRC) and not a human?	Yes: 7 / No: 0 / Unsure: 0 / Partly: 0	due to clear assignment of the flights, but a complete understanding not feasible in 2 days study; basic tasks well taught and introduced, automatic delegation seen critically (ATCOs' picture difficult to establish and refresh)
		Did you understand the role and limitations of the assistant during the test?	Yes: 4 / No: 2 / Unsure: 0 / Partly: 1 	

Figure 5. LOKI findings regarding EU AI Act, art 13 – Transparency.

Source **transparency** (meaning knowing whether the suggestion/action is coming from a human or from an AI system) worked generally well. Deeper role transparency (meaning fully understanding DIRC's role, limitations and what was its logic behind its delegation of flights) was more split. This is an important difference as deeper transparency is more relevant for answering accountability questions, such as: what are the system's main tasks? Why is it involved now? What are its limits and when should the human intervene?

AIA requirement	Brief description	Questions	User response (n=7)	User explanations
Art. 14 – Human Oversight	The system must allow human control and override (when needed)	Did you feel you were always in control of the system? Were you confident you could override or ignore the system if needed?	Yes: 4 / No: 2 / Unsure: 0 / Partly: 1 	weaknesses, capabilities and handling of DIRC absolutely clear; quickly learnt what didn't fit, delegation then cancelled; Manageable functional allocation, feeling of control lacked, because it is unclear why and when flights were assigned to DIRC (knowledge of assignment rules would be helpful here)
		Did the system's design (alerts/ interface) help you make your own decisions?	Yes: 7 / No: 0 / Unsure: 0 / Partly: 0	

Figure 6. LOKI findings regarding EU AI Act, art 14 - Human Oversight.

Human oversight was generally operationally possible as most ATCOs felt able to override or ignore DIRC and the interface itself also supported this decision-making task. However, some reported reduced feeling of control because it was unclear why and when some of the flights were assigned to DIRC instead of themselves. This is an important finding as oversight becomes meaningful only when the human understands the allocation rules and the limits the system has. This is the only way to be able to intervene when and where necessary.

AIA requirement	Brief description	Questions	User response (n=7)	User explanations
Art. 15 – Accuracy, Robustness & Cybersecurity	The system must perform reliably, be resilient and remain safe under attack/failure/error	Did DIRC provide reliable and accurate recommendations/ actions in different scenarios?	Yes: 2 / No: 1 / Unsure: 1 / Partly: 3 	basic DIRC tasks (sending off, taking over the planes until the trip) worked quiet well; some inappropriate clearances by DIRC (which were with the rules but were not safe); information on planned actions by DIRC was lacking for trusting DIRC
		Did you trust the system to remain safe under heavy traffic or stress?	Yes: 3 / No: 1 / Unsure: 0 / Partly: 3 	

Figure 7. LOKI findings regarding EU AI Act, art 15 - Accuracy, Robustness, Cybersecurity.

In general, the end users were cautiously positive about **reliability**, especially for routine tasks, but they appeared less confident in heavy traffic or stress edge cases. They also reported a couple of cases where DIRC's suggestion may have been rule compliant, but possibly not completely operationally safe enough. This shows that compliance and formal rules do not fully guarantee operational safety. Professional human judgement takes into consideration context-related

assessments that are difficult to reduce to the rule of compliance alone. This is why guidance must take into account operational realities.[30]

AIA requirement	Brief description	Questions	User response (n=7)	User explanations
Cross-cutting: Art. 4 – AI Literacy	People interacting with AI must have adequate training and literacy to use it safely and effectively	Did you receive enough training/ information to feel confident in using DIRC?	Yes: 7 / No: 0 / Unsure: 0 / Partly: 0	basic DIRC tasks (sending off, taking over the planes until the trip) worked quiet well; some inappropriate clearances by DIRC (which were with the rules but were not safe); information on planned actions by DIRC was lacking for trusting DIRC
		Did you feel you had the knowledge to judge when to rely on or to ignore DIRC's advice?	Yes: 7 / No: 0 / Unsure: 0 / Partly: 0	

Figure 8. LOKI findings regarding EU AI Act, art 4 - AI literacy.

AI acceptance	Brief description	Questions	User response (n=7)	User explanations
Cross-cutting: Acceptance **(not an AI Act requirement, but highly relevant)**	System adoption in real-life applications depends on trust, usability and willingness to integrate it into daily operations	Based on your experience, would you feel comfortable using DIRC in real ATC operations?	Yes: 7 / No: 0 / Unsure: 0 / Partly: 0	mainly step-by-step introduction desired (4/7) starting with routine tasks for DIRC, possibly initially with delegation by ATCOs
		Would you trust and accept DIRC as part of your daily work, if improvements are made?	Yes: 7 / No: 0 / Unsure: 0 / Partly: 0	

Figure 9. LOKI findings regarding general acceptance.

Two "horizontal" requirements outside the context of the AI Act's high-risk systems requirements were also considered for this study and they focused on **AI literacy and acceptance**. All ATCOs felt that they received sufficient training in order to properly use DIRC and they all seemed open to using DIRC in real life if improvements were made. They indicated a preference in a step-by-step introduction of such technologies, starting with routine tasks and flight delegation remaining under ATCO control. This suggests that hesitance is not necessarily connected to the AI system itself, but rather on uncontrolled or unexplained changes in authority between the system and the human.

Reviewing all these LOKI findings together, we can see that the accountability concerns that the end users have cannot be found solely in high-level regulation, as they tend to ask for more understandable roles, authority, limitations, explanations, oversight and deployment phases.

[30] The cybersecurity dimension of art 15 of the EU AI Act was not specifically assessed in this study, as it was deemed as less directly accessible to end-user perception and requires more technical assurance evidence, such as threat analyses or security tests. See AI Act, supra note 11, art 15.

These are exactly the kind of issues that sectoral guidance, such as the recently published EASA Concept Paper Issue 03, can address.

Section 5: Mapping of LOKI findings to EASA Concept Paper Issue 03

The EASA Concept Paper: Guidance for safety-related Artificial Intelligence applications, Proposed Issue 03[31], is an important document for understanding how the aviation sector is attempting to translate high-level AI governance requirements into sector-specific expectations. The document is not legally binding itself and it is also in its draft version, therefore it does not provide final means of compliance. However, it sets a structured framework of objectives and anticipated means of compliance that can guide organizations through the development, approval and oversight of AI applications that are used in an aviation safety-related context. Reading through the document, one can say that it falls under a middle regulatory space: it is far more concrete than the AI Act's high-level requirements, but less prescriptive than detailed certification and operational procedures, so it can enable adaptation per case.

The relevance of the Concept Paper 03 lies in the fact that it addresses this middle space between legal requirements and operational implementation which is central to this paper. The EASA Ethics for AI in Aviation survey[32] shows that aviation professionals asked for regulation on responsibility and accountability matters in human-AI teaming. The LOKI findings helped better understand what some of these needs may be in practice and these needs suggested that at least some of the arisen accountability concerns may not require new hard law, but rather more mature sector-specific guidance to be addressed. For this reason, the LOKI findings were mapped to the main areas (objectives, anticipated means of compliance and organizational provisions) discussed in the Concept Paper in order to see if there is any relevance and what this relevance can show us regarding responsibility and accountability matters. This mapping can be found in the table below.

LOKI finding	Concept Paper 03 relevant area	Accountability relevance and guidance
ATCOs generally perceived DIRC as safe, but there were isolated cases of possibly unsafe suggestions.	**CH-02/CH-03** High-level tasks, usage domain, operational scenarios, degraded scenarios and operational domain	Accountability cannot rely on nominal system behavior. Guidance should advise organizations to strongly consider edge cases and operational feedback as part of the AI system's risk management lifecycle.
	RA-01/RA-02 Functional decomposition, risk assessment per scenario	
	ORG-03/ORG-05 Continued risk assessment, adaptation of safety/risk management	
Users did not observe biased or systematically strange system behavior, but this cannot prove data quality or complete absence of bias.	**SU-DA-04/SU-DM-06/R-DA-04** Data/scenario quality, completeness and representativeness	Users could detect visible anomalies, but they cannot validate the entire data governance process. Guidance should clarify how user perception complements, but does not replace technical assurance.
	ET-01/ORG-04 Ethics-based assessment, continued ethics monitoring	
	ORG-05 Reporting potential risks/bias	

[31] EASA Concept Paper Issue 03, supra note 7.
[32] EASA Ethics Report, supra note 5.

LOKI finding	Concept Paper 03 relevant area	Accountability relevance and guidance
ATCOs has different views of "enough information" provided by DIRC in order to understand its recommendations or its plans for next steps.	**EXP-02/EXP-03** Characterizing the need for operational explainability, documenting explainability methods **SURK-EXP-01** Development and operational explainability	Accountability requires explanations that are meaningful to the operator, to the context and to the timing (not just available in technical documentation).
ATCOs valued logging and history, but the feature was not fully embedded in their routine workflow.	**CRA-02/CRA-03** Operational data and metrics for continuous risk assessment **ORG-03** Continued risk assessment based on operational data, in-service events, end-user evaluation **SURK-CM-01** Configuration management, archiving, retrieval, retention	Logs support accountability when they are available for reviews, learning or oversight. Guidance should portray logging as part of the operational learning and training and not just as a technical compliance requirement.
Source transparency was clear, but understanding DIRC's full role, limits and delegation of flights logic was more mixed.	**CH-01/CH-02/CH-03** End users, roles, responsibilities, high-level tasks, usage domain **CL-01/CL-02** Decision-making chains, end-user authority, classification of the AI-based system	Source transparency is not enough for accountability. Aviation professionals also need to understand what the AI system is expected to do, when, what authority it has and what its limits are.
Users felt able to override DIRC, but some reported reduced control because the assignment rules were a bit unclear.	**CL-01/CL-02** Authority allocation, AI level classification **EXP-02/EXP-03** Explainability needs for AI outputs **RS-01/RS-02** Responsibility scheme assessment, assigned responsibilities supported by authority-access-capability.	Meaningful oversight is not only about the possibility to override the AI system. Oversight should be informed by sufficient information on task allocation, authority and the conditions under which intervention is expected.
ATCOs were cautiously positive about reliability, especially for routine tasks. Confidence could be improved for heavy traffic/ stress edge case.	**CH-03** Operational degraded scenarios **RA-01/RA-02** Risk assessment per scenario **SA-03/SA-04/SA-05** Uncertainty, degraded values, exposure outside the OD/ODD **SU-LM-10/SU-LM-12** Performance verification, stability across ODD	This promotes a guidance-based approach that distinguishes routine from edge cases and encourages the organizations to justify the operational contexts in which such AI systems are planned to be introduced.
Users felt sufficiently trained to use DIRC within the context of this study, but some	**ORG-07/ORG-08** End user trainings based on the AI specificities **EXP-02/EXP-03/CL-01/CL-02**	Training is absolutely necessary, but it cannot compensate for an unclear design or responsibility

LOKI finding	Concept Paper 03 relevant area	Accountability relevance and guidance
aspects on system design or explanations could be improved.	Explanation and authority clarity (not just training) per AI level classification	allocation. Guidance should connect competences with system transparency and organizational responsibility.
ATCOs were open to use DIRC in daily operations if improvements were made and if a step-by-step introduction was followed.	**CL-02** AI level classification	This shows that acceptance can be built gradually. Guidance should support proportionate and phased introduction steps.
	CH-03 Limiting and defining usage domain	
	ORG-03 Continued monitoring through operational data and in-service events	
ATCOs wanted clearer responsibility and accountability allocation in human-AI teaming.	**CL-01/CL-02** Authority, AI classification	Responsibility allocation becomes the basis of accountability, especially when authority is shared or (partially) delegated.
	RS-01 Establishing and documenting the responsibility scheme	
	RS-02 Responsibilities through authority-access-capability	

Table 1. Mapping of LOKI findings to EASA's objectives and concepts introduced in its Concept Paper Issue 03.

This mapping shows three important points: (1) the LOKI findings are well in line with the structure of the Concept Paper, as the points that the ATCOs raised correspond to the areas that EASA is already trying to structure with this guidance material. This strengthens the arguments that the demand for more regulation might sometimes mean a demand for more mature guidance. (2) Accountability in aviation (or in any safety-critical domain) depends not only on technical assurance, but also on a set of organizational and operational conditions, such role-based responsibilities or usable explainability. The Concept Paper Issue 03 provides regulators and applicants/organizations with a common vocabulary and structure for also addressing these conditions. (3) The provided guidance should not be overly detailed as conditions change per scenario. The LOKI findings do not point at one missing rule or fixed technical solution, but rather to questions about role clarity, authority, explainability. This is why the Paper's guidance is useful, as it helps professionals follow the right analytical path without providing the same detailed solution for every aviation AI application.

Section 6: Discussion – From demands for regulation to better implementation guidance

Coming back to the original thesis of this paper, whether aviation professionals' demand for "more regulation" should be understood as a demand for additional hard law or as a demand for better implementation guidance, it is understood that the answer is not necessarily one versus the other. Binding, high-level rules remain necessary in order to establish legal obligations regarding responsibilities and safety requirements. However, the EASA Ethics Survey and the LOKI findings showed that many accountability-related concerns arise at the implementation level. These professionals do not only want to know that AI systems are regulated; they also want to understand how these systems will be introduced, how roles and tasks will be allocated between the system and the professional, how human oversight will remain meaningful and how responsibility (and naturally, accountability) will be distributed when humans and AI collaborate or cooperate.

This is where guidance becomes particularly important. The AI Act and other relevant regulatory texts provide high-level legal requirements, but they cannot fully explain how these should be translated into every aviation operational context. For example, they can require "transparency" or "human oversight", but they cannot define what kind of explanations an ATCO needs during a time-critical interaction with an AI-enabled assistant or how responsibility should be allocated when some tasks are partially allocated to an AI-based system. These questions need a clearer bridge between legal principles and operational practice. The LOKI findings indicatively show where this bridge is most needed, while EASA's Concept Paper Issue 03 shows how this bridge can begin to take form.

In this sense, the Concept Paper Issue 03 does not add an additional legal layer to the discussion, but rather it provides a common vocabulary and structure for dealing with the questions and patterns emerging in aviation scenarios, such as the ones that arose within the context of LOKI: defining usage domains, classifying AI levels (as also indicated in regulatory text – EU AI Act), assessing human-AI interactions, supporting operational explainability, linking responsibility with authority and applying the objectives in a proportionate manner. This is crucial as accountability in AI-based aviation systems requires the human who is expected to act responsibly to have the necessary authority, information, skills and their organization's support to do so – not just to have the responsibility formally.[33]

The first recommendation to regulators is therefore to treat such calls for "more regulation" as signals of governance uncertainty and not as regulatory gaps that need to be filled with more hard law. In some cases, new binding rules will be necessary, but in many others, the more appropriate response would be clearer guidance with better explanation of the existing obligations, domain specific examples that were extracted with the help of the directly involved aviation professionals as well as more structured acceptable means of compliance. Regulators should therefore follow EASA's footstep and try to distinguish among legal gaps, guidance gaps, assurance gaps or competence gaps before deciding on which response is the most fitting.

Second, responsibility and accountability should be addressed early in the AI application lifecycle and not only at the end of the approval or deployment phases. For human-AI teaming in particular, responsibility should be connected to the system's role and level of AI-enabled support; meaning, whether the system solely assists the human, cooperated with them, collaborates more closely or performs parts of the tasks. This assessment in combination with task allocation, human authority and the information available to the end user could help avoid responsibility gaps. It also mitigates the exact opposite problem of responsibility dumping: shifting the responsibility to human operators who formally remain "in the loop", but do not have enough information or control to exercise meaningful oversight.

Third, regulators should continue developing guidance on operational explainability treating is as a core element of accountability. The LOKI feedback shows that users may know when an output comes from an AI system (source transparency), but they would have liked to be able to understand more DIRC's roles, limits and next steps. Guidance should help applicants justify why the explanation approach they chose is appropriate for each operational scenario (meaning roles, timings, safety contexts, workload). At the same time, this guidance should remain principle-based and proportionate, rather than describing in detail the interface or the explanation format that should be used for each aviation AI system.

[33] See AI Act, supra note 11, art 14; Cavalcante Siebert et al, supra note 17.

Fourth, regulators should promote step-by-step introductions of AI-based systems that are used in the safety-critical contexts. The LOKI end users appeared generally very open to include DIRC to their everyday work life if improvements were made and if introduction happened gradually, starting with routine tasks first. Such an approach would help organizations and authorities to test whether the measures they adopted regarding human oversight, role allocations and explainability are actually valid in simpler realistic-like settings, before moving to full adoption and dealing with edge cases.

Finally, regulators should continue engaging with the directly-involved aviation professionals in the development and refinement of AI guidance. EASA has the third issue of the Concept Paper open for review and comments in order to collect as much feedback as possible from the front-line professionals that are dealing with such topics. This feedback cannot replace technical assurance or legal compliance evidence, but it can reveal where regulatory concepts are unclear in practice.

All in all, the findings suggest that better guidance should not be viewed as a weaker version of regulation. When read together – regulation and guidance – they become most useful and usable. The challenge for EASA and any other relevant regulator is to find the best balance between being concrete enough to guide applicants/organizations, but not too detailed so they become too inflexible or unsuitable for various scenarios.

Section 7: Conclusion

In this paper, we examined whether the aviation professionals' demand for "more regulation" on responsibility and accountability in aviation AI-based systems should necessarily be understood as a need for more hard law. Basing our research on the EASA Ethics for AI in Aviation survey, the LOKI ELSA assessment with ATCOs interacting with DIRC and EASA's latest draft version of the Concept Paper Issue 03, we argue that this demand might be interpreted differently. Aviation professionals do not simply ask for AI systems to be regulated, they also ask how such systems should be introduced, how roles and tasks should be allocated between the humans and the system, how human oversight can remain meaningful and how responsibility and accountability are distributed among the involved actors. The LOKI findings showed that the ATCOs were generally open to a gradual integration of such systems (especially for routine tasks), but their acceptance depended on clearer role allocation, better explainability and more understanding of the system's limits and logic.

The key takeaway is that in such settings accountability needs more than just assigning formal responsibilities. It requires responsible actions that come with sufficient authority, information that is usable and useful, adequate skills and competences, proper explanations and organizational support. The mapping of the LOKI findings to the EASA's Concept Paper Issue 03 shows that many of the topics the ATCOs raised are in line with the areas that EASA is already trying to structure. While much of the current discussion on AI governance focuses on legal obligations, technical assurance, standardization or explainability in general, this paper tried to connect the aviation professionals' demand for regulation with the concrete user feedback received from a human-AI teaming scenario within the context of the research project LOKI. The paper also showed that sector-specific guidance can act as a bridge between legal requirements and operational practice.

In this sense, guidance should not be viewed as weaker version of regulation, but as one of the main ways this regulation becomes usable in a complex, safety-critical domain, such as aviation. The analysis indicates that documents as EASA's Concept Paper 03 can help regulators,

organizations and end users to establish a common route towards an implementation process that is meaningful and feasible. Future research should build on this approach by applying similar mappings to other aviation scenarios and domains, larger user groups or more mature AI applications. Such work could support the development of aviation AI governance that is legally and technically robust, but also understandable, acceptable and operationally feasible for the professionals that are expected to work with these systems.

Bibliography

Legislation

Regulation (EU) 2018/1139 of the European Parliament and of the Council of 4 July 2018 on common rules in the field of civil aviation and establishing a European Union Aviation Safety Agency, and amending Regulations (EC) No 2111/2005, (EC) No 1008/2008, (EU) No 996/2010, (EU) No 376/2014 and Directives 2014/30/EU and 2014/53/EU, and repealing Regulations (EC) No 552/2004 and (EC) No 216/2008 and Council Regulation (EEC) No 3922/91, [2018] OJ L 212/1.

Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (AI Act) and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828, [2024] OJ L 2024/1689.

Secondary material: Reports

European Aviation High Level Group on Artificial Intelligence, *The FLY AI Report: Demystifying and Accelerating AI in Aviation/ATM* (Brussels: EUROCONTROL, 2020), online: <eurocontrol.int>.

European Union Aviation Safety Agency, *Artificial Intelligence Roadmap 2.0: A Human-Centric Approach to AI in Aviation* (Cologne: EASA, 2023), online: <easa.europa.eu>.

European Union Aviation Safety Agency, *EASA Artificial Intelligence Concept Paper — Proposed Issue 3* (Cologne: EASA, 2026), online: <easa.europa.eu>.

European Union Aviation Safety Agency, *Ethics for Artificial Intelligence in Aviation: Aviation Professionals Survey Results 2024/2025* (Cologne: EASA, 2025), online: <easa.europa.eu>.

Secondary material: Articles

Bainbridge, Lisanne, "Ironies of Automation" (1983) 19:6 Automatica 775, DOI: <10.1016/0005-1098(83)90046-8>.

Bovens, Mark, "Analysing and Assessing Accountability: A Conceptual Framework" (2007) 13:4 European Law Journal 447, DOI: <10.1111/j.1468-0386.2007.00378.x>.

Bryson, Joanna J, Mihailis E Diamantis & Thomas D Grant, "Of, For, and By the People: The Legal Lacuna of Synthetic Persons" (2017) 25:3 Artificial Intelligence and Law 273, DOI: <10.1007/s10506-017-9214-9>.

Busuioc, Madalina, "Accountable Artificial Intelligence: Holding Algorithms to Account" (2021) 81:5 Public Administration Review 825, DOI: <10.1111/puar.13293>.

Cavalcante Siebert, Luciano et al, "Meaningful Human Control: Actionable Properties for AI System Development" (2023) 3 AI and Ethics 241, DOI: <10.1007/s43681-022-00167-3>.

Elish, Madeleine Clare, "Moral Crumple Zones: Cautionary Tales in Human-Robot Interaction" (2019) 5 Engaging Science, Technology, and Society 40, DOI: <10.17351/ests2019.260>.

Endsley, Mica R, "Supporting Human-AI Teams: Transparency, Explainability, and Situation Awareness" (2023) 140 Computers in Human Behavior 107574, DOI: <10.1016/j.chb.2022.107574>.

Hauptman, Allyson I et al, “Understanding the Influence of AI Autonomy on AI Explainability Levels in Human-AI Teams Using a Mixed Methods Approach” (2024) 26 Cognition, Technology & Work 435, DOI: <10.1007/s10111-024-00765-7>.

Koivisto, Ida, Riikka Koulu & Stefan Larsson, “User Accounts: How Technological Concepts Permeate Public Law Through the EU’s AI Act” (2024) 31:3 Maastricht Journal of European and Comparative Law 412, DOI: <10.1177/1023263X241248469>.

Lozano Tafur, Cristian et al, “Applications of Artificial Intelligence in Air Operations: A Systematic Review” (2025) 25 Results in Engineering 103742, DOI: <10.1016/j.rineng.2024.103742>.

Nissenbaum, Helen, “Accountability in a Computerized Society” (1996) 2:1 Science and Engineering Ethics 25, DOI: <10.1007/BF02639315>.

Novelli, Claudio, Mariarosaria Taddeo & Luciano Floridi, “Accountability in Artificial Intelligence: What It Is and How It Works” (2024) 39:4 AI & Society 1871, DOI: <10.1007/s00146-023-01635-y>.

Roberts, Karlene H, “New Challenges in Organizational Research: High Reliability Organizations” (1989) 3:2 Industrial Crisis Quarterly 111, DOI: <10.1177/108602668900300202>.

Santoni de Sio, Filippo & Giulio Mecacci, “Four Responsibility Gaps with Artificial Intelligence: Why They Matter and How to Address Them” (2021) 34:4 Philosophy & Technology 1057, DOI: <10.1007/s13347-021-00450-x>.

Sukenik, Nili, “Coordination in the Cockpit: A Game-Theory View of Standard Operation Procedures” (1998) 8:4 The International Journal of Aviation Psychology 405, DOI: <10.1207/s15327108ijap0804_5>.

Wiegmann, David A et al, “Safety Culture: An Integrative Review” (2004) 14:2 The International Journal of Aviation Psychology 117, DOI: <10.1207/s15327108ijap1402_1>.

Woodlock, John & Håkan Hydén, “(f)Lex Avionica: How Soft Law Serves as an Instrumental Mediator Between Professional Norms and the Hard Law Regulation of European Civil Aviation Maintenance” (2020) 121 Safety Science 54, DOI: <10.1016/j.ssci.2019.08.037>.

Secondary material: Conference Papers

Dmitriev, Konstantin, Johann Schumann & Florian Holzapfel, “Toward Certification of Machine-Learning Systems for Low Criticality Airborne Applications” in *40th Digital Avionics Systems Conference, DASC 2021 — Proceedings* (New York: IEEE, 2021), DOI: <10.1109/DASC52595.2021.9594467>.

Raji, Inioluwa Deborah et al, “Closing the AI Accountability Gap: Defining an End-to-End Framework for Internal Algorithmic Auditing” in *FAT ’20: Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (New York: Association for Computing Machinery, 2020) 33, DOI: <10.1145/3351095.3372873>.

Secondary material: Websites

Deutsches Zentrum für Luft- und Raumfahrt, “LOKI – Kollaboration von Luftfahrt-OperateurInnen und KI-Systemen” (last visited 23 July 2026), online: <dlr.de>.