Advances in Human Factors of Transportation, Vol. 224, 2026, 8–17
https://doi.org/10.54941/ahfe1007828

The Role of Fatigue Risk Management Systems (FRMS) in the Implementation of Human–Artificial Intelligence (AI) Teaming in the Aviation Ecosystem

Debra Henneberry[1], Dimitrios Ziakkas[2,3], Eleftherios Bokas[1], and Anastasios Plioutsias[2]

[1]Purdue University, School of Aviation and Transportation Technology, West Lafayette, IN 47907, USA

[2]Coventry University, Faculty of Engineering, Environment and Computing, Coventry, CV1 5FB, U.K.

[3]Sky Express, Athens International Airport "El. Venizelos" - Building 23, Athens, 190 19, Greece

ABSTRACT

The emergence of Human-AI teaming within the aviation ecosystem introduces profound implications for fatigue management, operational resilience, and safety assurance. As artificial intelligence systems become increasingly embedded in flight operations, predictive analytics, crew scheduling, and cockpit decision support, the integration of Fatigue Risk Management Systems (FRMS) takes on an expanded and critical role. This paper examines how FRMS principles can support, and must be adapted to support, the safe and effective implementation of human-AI teaming, ensuring that technological innovation enhances rather than undermines human performance in safety-critical environments. The paper situates fatigue as a persistent human factor risk that continues to shape error pathways, cognitive performance, and decision-making under operational stress. AI-enabled systems hold the potential to augment fatigue management by providing real-time physiological monitoring, predictive fatigue modelling, adaptive workload distribution, and decision-support cues during periods of reduced alertness. However, the introduction of AI also reshapes the operational landscape: as tasks shift between human and machine, the cognitive workload profile of pilots may fluctuate unpredictably. In highly automated or single-pilot contexts, AI systems may inadvertently increase fatigue risks by imposing vigilance demands, imposing excessive monitoring burdens, or providing poorly calibrated levels of assistance. These emerging risks underscore the need for FRMS frameworks capable of recognising and governing the unique human-machine interactions associated with advanced automation. The paper examines how each core component of FRMS, policy, risk assessment, data-driven monitoring, and training, contributes to the governance of human-AI teaming. FRMS Policy must explicitly acknowledge AI as both a potential fatigue mitigator and a fatigue hazard, emphasising a human-centric philosophy that prioritises pilot wellbeing alongside operational efficiency. Fatigue Risk Assessment processes must incorporate new hazard categories, including over-reliance on algorithmic alerts, cognitive underload resulting from task redistribution, and increased monitoring pressures associated with supervising autonomous functions. Safety

Received February 6, 2026; Revised April 9, 2026; Accepted April 23, 2026; Available online July 20, 2026

Assurance within FRMS must transform into a continuous, adaptive process capable of monitoring AI performance, evaluating AI-human workload balance, and detecting early signs of fatigue-induced system drift. The paper proposes an integrated FRMS-AI governance model tailored to the future of human-AI teaming. This model incorporates predictive fatigue analytics, adaptive automation strategies, human-machine workload harmonisation, and AI-specific fatigue indicators within FRMS oversight. The findings emphasise that the success of human-AI teaming in aviation will depend not solely on technological progress but on the robustness of FRMS to manage evolving human cognitive vulnerabilities in increasingly automated operational ecosystems.

Keywords: Fatigue Risk Management Systems (FRMS), Human-AI teaming, Aviation fatigue, Adaptive automation, Cognitive workload, Predictive analytics

INTRODUCTION

Fatigue has long been recognised as a persistent and systemic contributor to degraded human performance in aviation, influencing attention, working memory, risk perception, and decision-making under time pressure. Despite decades of regulatory oversight, prescriptive flight and duty time limitations alone have proven insufficient to address the complex, context-dependent nature of fatigue in modern operations. The emergence of Fatigue Risk Management Systems (FRMS) represented a paradigmatic shift away from purely rule-based compliance toward performance-based, data-driven, and human-centred risk governance (ICAO, 2015; FAA, 2010).

In parallel, the aviation ecosystem is undergoing a profound transition as artificial intelligence is increasingly deployed across flight operations, scheduling, decision support, and predictive safety analytics. Human-AI teaming is no longer confined to experimental domains; it is becoming an operational reality that reshapes task allocation, cognitive workload distribution, and the nature of human oversight in safety-critical environments (Parasuraman & Riley, 1997; Hollnagel, 2014). While AI systems are frequently positioned as fatigue mitigators, capable of maintaining vigilance, detecting anomalies, and supporting decision-making, their integration introduces new, often under-examined fatigue pathways.

Fatigue in aviation cannot be reduced to physiological sleep debt alone. It is an emergent condition shaped by circadian disruption, task monotony, workload variability, organisational pressures, and the cognitive demands of monitoring complex systems. Human-AI teaming alters each of these dimensions. In highly automated environments, pilots may experience reduced manual engagement while simultaneously bearing increased responsibility for system supervision, anomaly detection, and intervention readiness. This combination of cognitive underload and sustained vigilance represents a well-documented precursor to fatigue-related performance degradation (Warm et al., 2008).

AI systems may inadvertently exacerbate these effects by smoothing operational variability, masking early warning cues, or reallocating tasks in ways that conflict with human cognitive rhythms. In single-pilot or reduced-crew concepts, the absence of social cross-checking further intensifies fatigue vulnerability, placing greater reliance on AI outputs during periods of reduced alertness (Ziakkas & Vink, 2023). From a human factors and safety science perspective, fatigue in human-AI systems manifests not only as slower reaction times or attentional lapses, but as impaired trust calibration, delayed challenge of automated decisions, and reduced capacity to recover from unexpected system states.

FRMS offers a structured mechanism to address these risks precisely because it treats fatigue as a system property rather than an individual failing. However, traditional FRMS implementations were designed for human-human crews and deterministic automation. The introduction of adaptive, learning AI necessitates a reconceptualisation of fatigue hazards to include cognitive mismatches between human and machine performance trajectories.

This paper argues that FRMS provides a uniquely suited governance framework for managing these emerging risks, but only if it is deliberately adapted to recognise AI not merely as a technical aid, but as an active participant in the socio-technical system. Without such adaptation, human-AI teaming risks reproducing historical patterns of automation surprise, vigilance decrement, and cognitive underload that have previously undermined safety in highly automated cockpits (Reason, 1997; Parasuraman et al., 2008).

FRMS COMPONENTS APPLIED TO HUMAN–AI TEAMING

The FRMS policy establishes the organisational philosophy for recognising, reporting, and managing fatigue. In the context of human-AI teaming, policy must explicitly acknowledge that AI systems may simultaneously function as fatigue mitigators and stressors. A human-centric FRMS policy reframes AI deployment as contingent on its demonstrated ability to support sustainable human performance rather than solely on operational efficiency. Such policy alignment is consistent with EASA's emphasis on human-centred automation and shared responsibility for safety outcomes across organisational levels (EASA, 2023). Importantly, FRMS policy must clarify accountability boundaries when AI-supported decisions intersect with fatigue-impaired human judgement, thereby avoiding ambiguity that can suppress reporting or foster inappropriate reliance on automation (EASA, 2023; Ziakkas & Chazapis, 2022).

Fatigue risk identification within human-AI systems must expand beyond duty hours and sleep opportunity to incorporate new hazard categories. These include prolonged monitoring of autonomous functions, task disengagement followed by sudden re-engagement demands, and reliance on algorithmic alerts during circadian low points. Cognitive underload, frequently underrepresented in traditional fatigue models, becomes a critical risk factor when AI absorbs routine tasks, leaving humans with episodic, high-stakes interventions. Table 1 illustrates how conventional fatigue hazards are transformed within human-AI teaming contexts and highlights the need for revised assessment criteria (Hollnagel, 2014).

Table 1: Fatigue hazard transformation in human-AI teaming.

Traditional FRMS Hazard	Human-AI Teaming Extension	Fatigue Implication
Long duty periods	Extended passive monitoring	Vigilance decrement
Night operations	AI-mediated task execution	Over-reliance during circadian lows
High workload phases	Sudden task handover from AI	Cognitive overload
Reduced rest	Continuous system supervision	Mental fatigue accumulation

Safety Assurance and Data-Driven Monitoring

Safety assurance represents the analytical core of FRMS, relying on continuous monitoring, trend analysis, and corrective action. In human-AI systems, assurance must encompass both human and machine performance indicators, recognising that fatigue-related degradation may manifest asymmetrically. AI systems may maintain nominal performance while human operators experience declining alertness, increasing the risk of unchallenged system errors.

Advanced FRMS implementations may leverage bio-mathematical fatigue models, physiological data streams, and operational performance metrics to detect early signs of human-AI misalignment. Such data must be interpreted within a Just Culture framework to ensure that fatigue signals associated with AI interaction are reported without fear of attribution or sanction (Reason, 1997). Table 2 outlines key safety assurance indicators relevant to human-AI teaming (Reason, 1997; Ziakkas & Plioutsias, 2024).

Table 2: FRMS safety assurance indicators for human-AI systems.

Indicator Category	Example Metric	Safety Relevance
Human performance	Reaction time variability	Fatigue detection
AI behaviour	Alert frequency trends	Over-automation risk
Interaction quality	Intervention latency	Trust calibration
System resilience	Recovery from AI errors	Fatigue-induced vulnerability

METHODOLOGY

This study adopts a structured, theory-informed methodological approach combining the ADDIE instructional design framework with the Research Onion model proposed by Saunders, Lewis, and Thornhill. This combined methodology enables a systematic examination of how Fatigue Risk Management Systems (FRMS) can be adapted to govern fatigue-related risks emerging from human-AI teaming in aviation operations. The approach is intentionally human-centred, reflecting the socio-technical nature of fatigue, automation, and organisational safety, and aligns with performance-based regulatory principles advocated by ICAO, FAA, and EASA.

Research Philosophy and Approach

At the outer layer of the Research Onion, the study is grounded in a pragmatic research philosophy, recognising that fatigue in human-AI systems cannot be fully understood through a single epistemological lens. Pragmatism allows integrating human factors theory, regulatory guidance, and operational practice, acknowledging that both subjective human experience and objective system performance data are essential for safety-critical decision-making. This stance is consistent with contemporary safety science, which views risk as an emergent property of complex socio-technical systems rather than a linear cause-and-effect relationship (Hollnagel, 2014).

The research follows a deductive-abductive hybrid approach. Existing theories of fatigue, automation, and human-machine interaction provide the initial analytical structure, while abductive reasoning supports the identification of novel fatigue pathways arising from AI-enabled operations. This is particularly relevant given the limited empirical precedent for sustained human-AI teaming in highly automated or single-pilot contexts.

Research Strategy and Design

In alignment with the Research Onion's strategy layer, the study employs a conceptual and analytical research design, drawing on regulatory documentation, peer-reviewed literature, and operational safety frameworks. Rather than empirical experimentation, which remains constrained in live aviation environments, the methodology focuses on systematic synthesis and model development. This approach is appropriate for emerging operational concepts where governance frameworks must precede large-scale deployment.

The ADDIE model is used as an organising structure within this strategy, enabling the translation of analytical findings into a coherent FRMS-AI governance framework. ADDIE is particularly suited to this study because FRMS itself is cyclical, adaptive, and learning-oriented, mirroring the iterative nature of instructional design and continuous safety assurance.

ADDIE Framework Application

Analysis Phase

The analysis phase involved a comprehensive examination of fatigue as a human-factors risk in advanced automation environments. Regulatory texts from ICAO, FAA, and EASA were reviewed alongside human factors literature addressing fatigue, vigilance, trust in automation, and cognitive workload. Particular attention was given to how AI-enabled task redistribution alters traditional fatigue pathways, including vigilance decrement, cognitive underload, and delayed readiness for intervention. This phase established the baseline requirements for adapting FRMS to human-AI teaming.

Design Phase

During the design phase, core FRMS components, policy, risk assessment, safety assurance, and safety promotion, were mapped against identified human-AI fatigue risks. This mapping process was informed by established FRMS guidance and human-centred design principles, ensuring alignment with performance-based safety management. Design decisions prioritised human wellbeing, trust calibration, and resilience, explicitly rejecting purely technology-driven optimisation strategies.

Development Phase

The development phase translated the conceptual design into an integrated FRMS-AI governance model. This included the formulation of AI-specific fatigue hazard categories, the identification of suitable safety performance indicators, and the alignment of training objectives with fatigue-related challenges in human-AI interaction. Tables and conceptual models were developed to support traceability between fatigue risks, FRMS controls, and operational outcomes.

Implementation Phase

Although full operational implementation lies beyond the scope of this paper, the methodology considers how the proposed framework could be embedded within existing Safety Management Systems (SMS). This includes integration with reporting systems, data monitoring processes, and training programmes. Implementation considerations were evaluated against regulatory expectations for non-punitive reporting and continuous improvement.

Evaluation Phase

The evaluation phase focuses on the theoretical validation of the proposed framework through consistency checks against established fatigue science, automation research, and regulatory principles. The cyclical nature of FRMS is emphasised, with evaluation positioned as an ongoing process that supports adaptation as AI capabilities and operational concepts evolve. This reflects the recognition that both fatigue and human-AI interaction are dynamic phenomena requiring continuous oversight.

Time Horizon and Data Sources

Consistent with the Research Onion's time horizon layer, the study adopts a cross-sectional analytical perspective, synthesising contemporary regulatory and academic sources. Data sources include international regulatory guidance, peer-reviewed human factors research, and authoritative texts on fatigue and safety management. No primary human subject data were collected, avoiding ethical constraints while ensuring methodological robustness.

Methodological Rigor and Limitations

Methodological rigor is supported through triangulation across regulatory, theoretical, and operational perspectives. However, the absence of empirical operational data represents a limitation, reflecting the current early stage of large-scale human-AI teaming implementation in aviation (Ziakkas & Chazapis, 2022). This limitation is acknowledged as an inherent characteristic of forward-looking safety research and reinforces the value of conceptual frameworks in shaping future empirical inquiry.

FINDINGS

Training and safety promotion constitute a highly visible operational dimension of Fatigue Risk Management Systems, yet they are also the components most sensitive to the introduction of artificial intelligence into aviation operations. Within the methodological framework adopted in this study, training is not treated as a standalone educational activity, but as a system-level intervention derived from the analytical and design phases of the ADDIE model and grounded in the pragmatic, socio-technical assumptions of the Research Onion (Parasuraman et al., 2008).

From a methodological standpoint, safety promotion represents the point at which abstract fatigue risk analyses are translated into cognitive, behavioural, and cultural competencies. In human-AI teaming contexts, this translation becomes particularly critical, as fatigue no longer manifests solely through physiological depletion but increasingly through degraded monitoring capability, trust miscalibration, and impaired sensemaking when interacting with intelligent systems. The methodology therefore, positions training as a dynamic safety control rather than a static compliance requirement.

Methodological Foundations of Human-AI Fatigue Training

During the "Analysis phase" of the ADDIE framework, fatigue-related hazards associated with human-AI interaction were identified through synthesis of human factors literature and regulatory guidance. These hazards include vigilance decrement during prolonged automation supervision, delayed recognition of AI system degradation under fatigue, and cognitive inertia when transitioning from passive monitoring to active control. The Research Onion's pragmatic philosophy underpins this analysis by recognising that both measurable indicators and subjective human experience influence fatigue perception and performance degradation.

The "Design phase" operationalised these findings by defining training objectives aligned with FRMS safety outcomes rather than procedural knowledge acquisition. Training content is therefore designed to enhance metacognitive awareness, enabling operators to recognise how fatigue alters their interaction with AI systems and to anticipate situations in which automation may amplify, rather than mitigate, human vulnerability. This design logic reflects established safety science principles that emphasise adaptive capacity and informed decision-making over rote adherence to procedures (Hollnagel, 2014).

Development of Human-Centred Training Content

In the "Development phase", training modules were conceptually structured to reflect real operational scenarios in which fatigue and AI interaction intersect. Rather than isolating fatigue education from automation training, the methodology integrates these domains, reinforcing the understanding that fatigue is a contextual and relational phenomenon. Crucially, the methodology emphasises that training must address not only system functionality but also system fallibility. Operators are encouraged to develop an informed scepticism toward AI outputs during fatigue-prone phases, without fostering distrust or rejection of automation. This balance aligns with empirical research on trust in automation, which demonstrates that both over-trust and under-trust are associated with increased safety risk (Parasuraman et al., 2008).

Implementation Within FRMS and SMS Structures

The "Implementation phase" situates training and safety promotion within existing FRMS and SMS infrastructures. Methodologically, this ensures that training is not an isolated intervention but part of an integrated safety governance system. Training outcomes are linked to fatigue reporting practices, safety performance indicators, and safety assurance feedback loops, reinforcing the cyclical nature of FRMS.

Within this framework, safety promotion also encompasses organisational communication strategies that normalise discussions of fatigue in human-AI contexts. The methodology explicitly supports a Just Culture environment, recognising that reluctance to report fatigue-related AI interaction issues can undermine the effectiveness of both FRMS and emerging automation strategies (Reason, 1997). Training, therefore, includes structured reflection on near-misses, automation surprises, and fatigue-related discrepancies between human and machine judgment.

Evaluation and Continuous Learning

The "Evaluation phase" of ADDIE is methodologically critical in ensuring that training remains responsive to evolving AI capabilities and operational realities. Rather than relying solely on knowledge assessments, the evaluation focuses on behavioural indicators, reporting trends, and changes in fatigue-related event patterns. This aligns with the Research Onion's emphasis on methodological coherence, ensuring that evaluation methods remain consistent with the study's pragmatic philosophy and analytical aims.

Feedback from training evaluation informs iterative refinement of both training content and FRMS controls, reinforcing the system's adaptive capacity. As AI systems evolve through updates and learning algorithms, the methodology ensures that training evolves in parallel, preventing the gradual erosion of human competence and situational awareness that has historically accompanied the expansion of automation.

DISCUSSION

The integration of human-AI teaming into aviation operations challenges existing assumptions about fatigue mitigation. While AI offers unprecedented opportunities for predictive analytics and workload optimisation, it also introduces subtle cognitive demands that traditional fatigue frameworks may overlook. FRMS, when properly adapted, provides a governance structure capable of balancing innovation with human limitations. Beyond individual training, the methodology frames safety promotion as an indicator of organisational readiness for human-AI teaming (Hollnagel, 2014; Ziakkas & Plioutsias, 2024). A mature FRMS-AI integration is characterised by open dialogue about fatigue, transparent communication about AI limitations, and leadership engagement in promoting human wellbeing as a safety priority. This reflects a broader human-centred safety philosophy underpinning the study, in which technology is viewed as a contributor to resilience only when embedded within a supportive organisational culture. Methodologically, this perspective reinforces the conclusion that training and safety promotion are not downstream activities but central mechanisms through which FRMS governs the evolving relationship between humans and intelligent systems. Their effectiveness depends on alignment with analytical insights, regulatory expectations, and continuous evaluation, ensuring that human-AI teaming enhances, rather than compromises, fatigue resilience in aviation operations.

This paper proposes that effective human-AI integration depends not on replacing human judgment, but on designing systems that respect cognitive rhythms, preserve meaningful engagement, and support recovery from fatigue-related degradation. Such an approach aligns with broader human-centred safety philosophies and reinforces the role of FRMS as a living system rather than a compliance artifact.

CONCLUSION

Human-AI teaming represents a transformative shift in aviation operations, with far-reaching implications for fatigue management and safety assurance. This paper demonstrated that FRMS offers a robust and adaptable framework for governing these changes, provided it evolves to recognise AI as an active participant in the fatigue risk landscape. By integrating AI-specific hazards into policy, risk assessment, safety assurance, and training, FRMS can safeguard human performance in increasingly automated environments. Ultimately, the success of human-AI teaming will depend less on technological sophistication than on the aviation system's ability to manage human vulnerability with transparency, scientific rigour, and respect for human cognitive limitations. FRMS stands as a critical enabler of this balance, ensuring that advances in automation enhance rather than erode the resilience of the human operators upon whom safety ultimately depends.

ACKNOWLEDGMENT

The authors thank Coventry University, Purdue University faculty members, Sky Express, and HF Horizons analysts for their invaluable feedback, which contributed to this work.

REFERENCES

EASA. (2023). *Human-centred automation in aviation: Principles, challenges and regulatory considerations*. European Union Aviation Safety Agency.

FAA. (2010). *Fatigue risk management systems for aviation safety* (Advisory Circular AC 120-103A). Federal Aviation Administration.

Hollnagel, E. (2014). *Safety-I and Safety-II: The past and future of safety management*. Ashgate.

ICAO. (2015). *Fatigue risk management systems (FRMS): Manual for regulators* (Doc 9966). International Civil Aviation Organization.

Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. *Human Factors, 39*(2), 230–253. https://doi.org/10.1518/001872097778543886

Parasuraman, R., Sheridan, T. B., & Wickens, C. D. (2008). Situation awareness, mental workload, and trust in automation: Viable, empirically supported cognitive engineering constructs. *Human Factors, 50*(3), 433–444. https://doi.org/10.1518/001872008X288394

Reason, J. (1997). *Managing the risks of organizational accidents*. Ashgate.

Saunders, M., Lewis, P., & Thornhill, A. (2019). *Research methods for business students* (8th ed.). Pearson Education.

Warm, J. S., Parasuraman, R., & Matthews, G. (2008). Vigilance requires hard mental work and is stressful. *Human Factors, 50*(3), 433–441. https://doi.org/10.1518/001872008X312152

Ziakkas, D., & Vink, L. S. (2023). Human-centred artificial intelligence in aviation operations. In E. Hollnagel, J. Woltjer, & R. Amalberti (Eds.), *Human-intelligent systems integration* (pp. xx-xx). CRC Press.

Ziakkas, D., & Chazapis, S. (2022). Human-centred safety management in complex socio-technical systems. CRC Press.

Ziakkas, D., & Plioutsias, A. (2024). Cultural ergonomics and safety governance in high-reliability aviation systems. In Advances in Human Factors and Ergonomics. Springer.