Cognitive Resilience and Human-AI Teaming in Air Traffic Control

Silvia Torsi[1,*], Stefano Bonelli[1], Anna Giulia Vicario[1], Alfonso Levantesi[1] and Hossein Mapar[1]

[1]*Deep Blue, Via Daniele Manin 53, 00185 Rome, Italy*

Abstract

A resilient system is capable of absorbing shocks, adapting, and reorganizing while maintaining its function: this implies a structure with feedback capabilities, self-regulation, and continuous learning. Designing resilience in Human-AI Teaming for air traffic control (ATC) means creating hybrid cognitive ecologies, where technology enhances human cognitive abilities through conscious co-evolution. Human-AI teaming can be conceived as real-time collaboration for conflict detection, trajectory management, and response to unforeseen events. AI can continuously monitor the airspace, anticipate anomalies, and suggest corrective actions, while the human provides contextual judgment, creativity, and the ability to manage ambiguity. Resilience thus emerges from the adaptive interaction between these two cognitive agents, creating a system that is more than the sum of its parts. In this context, ATC becomes a paradigmatic environment for testing a new epistemological alliance between natural and artificial intelligence, where resilience is not just a response to emergencies, but a continuous operational practice based on shared awareness, adaptability, and incremental learning.

Keywords

air traffic control, resilience, hybrid cognitive ecologies, abduction

1. Introduction

The Air Traffic Control (ATC) system exemplifies one of contemporary society's most complex and safety-critical domains. Reactive approaches to safety in ATC often involve adding elements to the system as corrective and anticipatory measures in response to potential incidents. Anyway, safety is not a stable asset, but rather a dynamic non-event, and the path to safety lies in continuously identifying the system's changing vulnerabilities [1]. Although commercial air transport remains one of the safest sectors, with an extremely low accident rate across millions of annual flights, the very nature of the domain makes every single failure catastrophic. The steady growth of global air traffic, estimated with a significant Compound Annual Growth Rate (CAGR) in pre-pandemic and recovering industry studies, exponentially increases the complexity and density of airspace, severely testing existing safety paradigms. This paper, therefore, proposes a conceptual framework for integrating human and artificial cognitive resilience, positioning itself as a perspective analysis. In this context, the integration of artificial intelligence (AI) opens up new scenarios, in which the concept of resilience plays a key role. Resilience is understood here not only as the system's ability to absorb shocks, but also as its capacity to monitor, learn, and adapt in real time. This paper seeks to explore how the resilience of human cognition and that of AI can act synergistically to support an effective ongoing process that ensures safety in ATC.

2. The Dangers of a Static and Additive Approach to Safety

While robustness implies resistance to known or expected errors, resilience involves active adaptability to unforeseen events or gradual degradation [1]. In the aviation domain, a static and corrective-based

HAIC 2025 – First International Workshop on Human-AI Collaborative Systems, editors Michele Braccini, Allegra De Filippo, Michela Milano, Alessandro Saffiotti, Mauro Vallati; October 25, 2025, Bologna, Italy

*Corresponding author.
silvia.torsi@dblue.it (S. Torsi)

approach may lead to poor awareness of system weaknesses, generating a false sense of stability and inducing a progressive decline in vigilance, a phenomenon known as "drift into failure" [2]. Two emblematic examples [3] are the accidents of Air France Flight 447 (2009) and the Space Shuttle Columbia disaster (2003). Flight AF447, which crashed into the Atlantic Ocean with 228 fatalities, followed a period during which no major incidents had occurred in European commercial aviation. This low perception of risk contributed to the underestimation of critical issues related to the icing of airspeed sensors (pitot tubes), and revealed a deficiency in crews' ability to respond to situations involving autopilot disconnection. Similarly, the Columbia disaster occurred after 17 years of shuttle missions without human loss since the Challenger tragedy in 1986. Those repeated successes fostered an environment of overconfidence and normalization of risk, leading NASA to ignore alarming technical signals regarding foam insulation impact on the shuttle wing. In both cases, the absence of recent incidents was not a sign of robustness, but rather of resilience that had not been actively cultivated—with fatal consequences. This paper focuses in particular on the systemic aspects of human reasoning and the computational capabilities of AI.

3. Resilience in Complex Systems

Throughout the 20th century, science witnessed a fundamental shift from a mechanistic and reductionist paradigm to a systemic and holistic one. Fritjof Capra [4] identifies in complexity sciences and systems biology the origin of a new way of thinking, in which phenomena are no longer explained by isolating their parts but by understanding them as interconnected nodes within dynamic networks.

This systemic view, inspired by cybernetics and ecology, emphasizes that emergent knowledge is based on relationships, patterns of organization, and self-regulating processes. What makes a system a "system" is the structure of interactions among its components: feedback loops, thresholds of adaptation, leverage points [5]. Systems possess the ability to maintain their identity even in the face of external disruptions. Therefore, system resilience is an emergent function of interconnection and the adaptive capacity embedded in the network of processes and actors. Meadows [5] stresses that a resilient system is one that can "bounce back" to equilibrium after a disturbance, but even more importantly, it is one that can learn and reorganize by itself. Resilience is thus a form of "evolutionary fitness" that encompasses the ability to adapt to changing contexts and anticipate potential vulnerabilities through continuous informational feedback. In parallel, Capra [4] offers a vision of living systems as dynamic networks capable of self-renewal and transformation: their stability derives from structural plasticity—the ability to change configuration while preserving overall integrity and identity. In this sense, a system is inherently dynamic and capable of learning, responding, and evolving.

Applying this perspective to the domain of Air Traffic Control (ATC) means recognizing the need for organizational and cognitive architectures that can anticipate, absorb, and adapt to sudden changes by integrating human expertise and computational capabilities in a synergistic way. In fact, in ATC this vision is crucial to understanding that safety requires an ongoing process of monitoring and adaptation. A resilient ATC system is one that can maintain safe performance even under unforeseen conditions, integrating operational disturbances—such as traffic surges, human errors, or technological malfunctions—without collapsing.

This implies the presence of distributed feedback mechanisms, flexible coordination between human operators and intelligent systems, and an organizational culture focused on systemic learning and prevention, rather than solely on reactive incident response.

3.1. Human Cognition and Resilience

The resilience of human cognition manifests in the ability to respond effectively to novel and critical situations through a combination of flexible, intuitive, and reflective decision-making processes. As described by Kahneman [6], human thinking operates through two distinct systems: System 1, which is fast, automatic, and intuitive, and System 2, which is slower, analytical, and deliberative. Abduction [7], a form of hypothetical inference that enables the generation of plausible explanations under uncertainty, plays a crucial role in human problem-solving, especially in dynamic and unpredictable environments like air traffic control. In high-pressure, high-complexity domains such as aviation, cognitive resilience does not rely on the dominance of one system over the other, but rather on their dynamic integration. The ability to rapidly switch between intuitive responses and analytical reasoning is a hallmark of expert problem-solving, and, in such contexts, intuition is not merely an automatic reaction, but a refined expression of tacit knowledge, developed through years of practice and consolidated into flexible cognitive schemas [8, 9]. These elements make the human mind a quintessentially resilient agent, capable of swiftly reconfiguring itself in the face of the unexpected.

An example of this is the emergency landing of US Airways Flight 1549 on the Hudson River in 2009 [3], where Captain Chesley Sullenberger adeptly diverged from conventional protocols, drawing upon a type of situation judgment grounded in expert intuition, swift sensemaking, and the management of distributed cognitive load. Within minutes, the crew integrated the immediate recognition of damage (System 1) with a rational evaluation of alternative options (System 2), demonstrating how cognitive resilience is grounded in embodied experience, evolved situation awareness, abductive inferential capacity and awareness/deployment of the context. Within human-AI teaming, this potential can be amplified by artificial intelligence, which can provide predictive and diagnostic support without displacing the creative, inferential, and adaptive role of human thinking.

3.2. Resilience in AI

Resilience in artificial intelligence refers to the engineered capacity of systems to preserve operational integrity and adapt effectively when confronted with unforeseen conditions that go beyond their original design parameters. In this sense, resilience is best understood as the ability to sustain reliable performance in dynamic and uncertain environments. Methodologically, resilience builds upon the foundations of machine learning (ML) and deep learning (DL), which enable systems to detect patterns and make decisions through exposure to large volumes of data [10]. In classical ML, resilience is commonly achieved via probabilistic approaches, such as Bayesian models, which estimate distributions of possible outcomes and continuously update beliefs when new information arises.

By contrast, resilience in DL derives from neural network architectures inspired by biological systems, where knowledge is distributed across many interconnected nodes. This representation creates redundancy, ensuring tolerance to noise and partial failure and allowing systems to degrade gracefully. In this way, ML fosters resilience through explicit uncertainty management, while DL achieves it through redundancy and hierarchical feature learning—together offering two complementary paradigms for adaptive and robust systems [10]. Such engineered resilience is pivotal for human–machine interaction, as it underpins the notion of joint cognitive systems [1].

Within this model, AI augments human cognition by processing vast streams of data to detect weak signals and anticipate risks, including the gradual drift toward the boundaries of safe operation—that is, the tendency of systems and organizations to move imperceptibly closer to safety limits under pressures for efficiency or resource constraints, until an unforeseen event pushes them beyond [11]. A critical dimension of resilient AI is its ability to evaluate its own reliability and communicate transparently. This requires the system to qualify its outputs by expressing confidence levels and providing intelligible explanations in line with explainability principles [12]. Transparency, in turn, is essential to cultivating calibrated trust, enabling human operators to discern when reliance on the AI is appropriate and when expert judgment should prevail [13]. Ultimately, embedding resilience into AI must be conceived as a human-centered design endeavour [14], requiring systems that foster mutual intelligibility and shared situation awareness in high-stakes environments.

3.3. Cognitive Complementarity

These characteristics make ML and DL particularly well-suited for applications in safety-critical domains such as Air Traffic Control (ATC), where the ability to operate in partially observable, dynamic, and highly variable environments is essential. In such scenarios, AI resilience not only supports operational continuity but enables adaptive co-evolution within human-AI teaming, extending the diagnostic, predictive, and decision-making capabilities of the entire sociotechnical system. Thus, algorithmic resilience can support the evolution of safety from a reactive function to a systemic and predictive process. These capabilities complement human cognitive resilience, supporting operators in maintaining situation awareness even under overload or ambiguity [15]. Moreover, the adoption of hybrid human-AI architectures, which leverage the inferential flexibility of humans and the adaptive scalability of AI, allows for the development of ATC systems that not only respond to unforeseen events but learn from them with a view toward continuous improvement.

The resilience emerging from human–AI teaming is maximized when both components — the human and the artificial intelligence — interact according to a systemic and distributed inferential logic. Humans are particularly skilled in abductive reasoning, ambiguity management, and reasoning with partial or unstructured information—cognitive traits that allow them to generate plausible hypotheses under uncertainty and rapid change [7]. AI, particularly through machine learning, excels in massive data processing, hidden pattern detection, and probabilistic forecasting based on large-scale trained models. When these two forms of intelligence are fused within a collaborative system, resilience is the result of the dynamic complementarity of cognitive capabilities [1].

Resilience thus becomes an emergent property of the interaction between human and artificial agents within a distributed, self-monitoring system that is structurally capable of adapting to uncertainty—an essential condition for long-term safety in contemporary airspace.

3.4. Challenges and Open Questions

Achieving resilient human-AI synergy requires navigating the complex challenge of trust calibration, a delicate balance that must avoid the pitfalls of both overtrust and undertrust [16].

On one hand, overtrust manifests as automation bias—an uncritical acceptance of AI suggestions that reduces operator vigilance and can lead to long-term skill degradation. On the other hand, undertrust causes operators to dismiss valid AI insights due to the system's opacity, undermining the very purpose of the collaboration. Bridging this divide between human understanding and machine reasoning is the primary role of Explainable AI (XAI) [17], which serves as a critical enabler for a resilient joint cognitive system. By providing understandable justifications for its outputs, XAI fosters the shared situation awareness [18] and calibrated trust [19] that are the hallmarks of a true cognitive partnership. This calibration is especially critical because the AI's own resilience cannot be taken for granted; its models remain vulnerable when confronted with situations outside their training distribution.

Therefore, the ultimate goal is to cultivate a transparent team where mutual monitoring, enabled by explainability, allows each agent to be aware of the other's limitations, creating true cognitive synergy.

4. Possible Models of Human-AI Teaming

Translating these concepts into the ATC domain, resilience becomes the system's ability to respond to unexpected events without compromising safety. In aviation, the collaboration between humans and AI—human-AI teaming—can be conceived as a distributed inferential process that extends the diagnostic, predictive, and responsive capabilities of the human-machine team. For example, AI-based ATC systems can anticipate airspace saturation, detect irregular traffic patterns, or identify early signs of human error through temporal and semantic analyses of operational data. According to Klein et al. [16] and Hutchins [20], an effective human-agent team requires mutual predictability, reciprocal directability, and shared situation awareness.

In aviation—and particularly in air traffic control—this translates into the creation of a distributed cognitive cycle connecting perception, interpretation, action, and learning in both the human and the artificial system. For instance, a controller may detect weak signals of a potential conflict through situation experience and expert intuition, while the AI simultaneously analyzes trajectories in real time at a systemic scale to uncover patterns invisible to the human eye.

This type of adaptive and collaborative monitoring reflects the logic of Dekker and Pruchnicki [2] in which safety is preserved through continuous adaptation to perturbations and minor deviations rather than strict adherence to procedures. In this model, humans and algorithms operate as intelligent nodes in a distributed inferential network. Their collaboration manifests in three key capabilities, supported by concrete examples:

1. Early detection of drift signals: In line with Dekker's theory of "drift into failure" [2], an AI can monitor vast operational datasets (e.g., communications, trajectory deviations, response times) across entire sectors and over long periods. It could thus identify a slow, progressive normalization of non-standard or riskier procedures—a phenomenon nearly invisible to a single operator focused on the tactical present.
2. Redundant yet functionally complementary decision-making: Faced with a potential conflict, the AI might propose an optimal solution based on efficiency and fuel consumption. The human controller, however, drawing on contextual knowledge (e.g., unmodeled predicted turbulence, military activity in an adjacent sector), could discard that solution in favor of a wider but strategically safer maneuver. The final decision emerges from the synthesis of these two perspectives.
3. Co-construction of shared explanations through transparent interactions: Embracing the principles of Explainable AI (XAI) [21], the system does not merely issue a command. Instead of suggesting "Reroute flight X," it would communicate: "Suggestion: Reroute flight X via waypoint Z. Reason: High probability of conflict (92%) with flight Y in 7 minutes due to unforecasted high-altitude winds. This route reduces conflict probability to <1% with a 3% additional fuel cost." This transparency is essential for building calibrated trust and truly shared situation awareness.

This inferential continuity, enabled by the cognitive synergy between humans and AI, constitutes a new form of generative resilience: capable of learning from events, redefining operational rules, and progressively increasing the system's fitness in relation to contextual variability.

5. Transparent Interfaces and Perceived Reliability

The critical point in achieving synergy between human systemic thinking and the systemic approach of artificial intelligence lies in the construction of shared situation awareness. For human–AI teaming to operate effectively in high-stakes environments such as Air Traffic Control (ATC), it is essential that both human and machine agents are able to comprehend and anticipate each other's actions, sustaining a transparent and interpretable flow of information. Ultimately, the resilience of the ATC system as a distributed cognitive system depends on the convergence of the operator's intuitive and rational inferences with the predictive and analytical capabilities of the AI system.

This requires the rationale behind decisions to remain accessible at all times, promoting a dynamic, situated, and reflective trust model. The growing interest in Explainable AI (XAI) technologies directly addresses this need: to provide human users with understandable justifications for system decisions, thereby supporting human inferential reasoning and reinforcing shared situation awareness [21]. Consequently, Explainable AI constitutes an enabling condition for the development of trust [19], defined as a delicate balance that allows human operators to avoid both overtrust and unwarranted undertrust toward the automated system. Pioneering studies [16] have shown that poorly calibrated trust is one of the primary causes of human error in complex, technology-assisted systems.

In the ATM domain, this translates into the design of interfaces where AI operates not as an opaque tool, but as a collaborative teammate, supporting tasks such as conflict detection and conflict resolution, while preserving the operator's strategic oversight and authority to intervene in unforeseen situations. This symbiotic arrangement lies at the core of next-generation European and American air traffic management programs, such as SESAR (Single European Sky ATM Research) [22] and NASA NextGen, where the human role is redefined as that of an adaptive supervisor, integrating human insights and algorithmic recommendations within a multilayered cognitive cycle [16].

6. Case Study: CODA, a Paradigm of Resilient Human-AI Teaming in Air Traffic Control

The CODA system [23] represents a hybrid human–machine team framework, wherein air traffic controllers (ATCOs) and AI-driven automation collaboratively execute ATC tasks. This cooperation dynamically adapts to the cognitive and operational state of the ATCO by employing progressive automation, AI-enabled decision-support tools, based on continuous monitoring of real-time and anticipated operator status. CODA operates by continuously assimilating multiple streams of data, integrating air traffic parameters with neurophysiological indicators of the ATCO's cognitive state—encompassing workload, attention, stress, fatigue, and vigilance. Through predictive modeling, the system forecasts future task demands and anticipates the corresponding cognitive conditions of the controller. This holistic, context-aware assessment allows CODA to intelligently manage and redistribute ATC responsibilities according to the operator's mental workload and environmental complexity.

The CODA workflow can be conceptualized as a continuous loop:

1. Data Acquisition: The system simultaneously collects traffic data and neurophysiological signals;
2. Predictive Modeling: The AI analyzes these streams to predict the controller's future workload;
3. Adaptive Task Allocation: Based on the prediction, the system proactively assumes or offloads specific tasks;
4. Transparent Interface: The AI's decisions and system status are clearly communicated to the operator;
5. Human Action and Feedback: The operator acts, and their response (both operational and neurophysiological) serves as new input for the cycle.

Although the project is still under development, preliminary simulations and human-in-the-loop tests have yielded encouraging results. These initial studies indicate a statistically significant reduction in controllers' perceived workload during high-density traffic scenarios and an improvement in conflict detection times compared to baseline conditions without CODA's support. The system can alleviate specific duties from the human operator, such as maintaining aircraft separation, preventing collisions, ensuring efficient and orderly traffic flow.

At the core of CODA's architecture is an advanced adaptive automation mechanism that dynamically modulates task allocation between human and machine. This is informed by continuous evaluations of cognitive load and neurophysiological signals, ensuring that automation supports the ATCO without supplanting their central role in decision-making. The adaptive automation strategy balances workload to prevent operator overload or underload, thus maintaining optimal vigilance and performance. In doing so, CODA embodies a cognitive teaming approach in which human and AI systems share operational goals, mutually monitor status, and continuously adapt to evolving demands. Shared situation awareness is foundational to CODA's effectiveness. The system's dynamic, interactive visualization interface integrates diverse data inputs into a coherent, real-time representation of the operational environment, encompassing both air traffic status and human cognitive metrics. This interface functions as a critical cognitive nexus, enabling transparent, bidirectional communication between human controllers and AI automation. By rendering system status, operational priorities, and constraints in an accessible and interpretable manner, the interface supports the development of a shared mental model. This mutual understanding facilitates coordinated decision-making, enhances predictability of AI behavior, and fosters calibrated trust—a balance of appropriate reliance and skepticism essential for effective human–AI collaboration in safety-critical environments.

Transparency is thus a pivotal design principle within CODA, enabling operators to anticipate, understand, and predict the AI's actions and recommendations. This transparency underpins a distributed cognitive system framework, wherein human and artificial agents collectively process, share, and act upon knowledge. Such a distributed cognitive ecology is crucial for resilient and adaptive responses to the dynamic, uncertain conditions characteristic of modern ATC operations. Drawing on the concept of cognitive resilience [1], CODA enables the socio-technical ATC system to anticipate, absorb, and adapt to unexpected changes, operational stressors, and suboptimal human performance states while maintaining safety and operational continuity. The system's neurophysiologically informed adaptive automation embodies anticipatory resilience, proactively redistributing tasks before the ATCO reaches critical fatigue, overload, or inattention thresholds.

Beyond technological innovation, the CODA project highlights that integrating AI into air traffic control represents a profound epistemic and cognitive transformation. Resilience emerges through synergistic collaboration and continuous co-adaptation between human expertise and AI augmentation. In this respect, CODA exemplifies the theoretical framework of distributed cognition [20], whereby AI functions as a cognitive extension of the human operator, augmenting inferential processes, situation awareness, and adaptive decision-making.

Ultimately, CODA envisions a new ontology for air traffic control—one where knowledge, agency, and responsibility are dynamically distributed across a hybrid cognitive system composed of human and non-human actors. This paradigm shift redefines control as an emergent property of integrated human–AI collaboration, fostering a resilient, adaptive socio-technical ecosystem capable of addressing the increasing complexity and demands of future airspace operations.

7. Conclusions

This paper has presented a conceptual framework for analyzing resilience in Human-AI Teaming systems for air traffic control. It has been argued that in safety-critical domains such as aviation, the integration of artificial intelligence and human operators gives rise to a meta-system in which the inferential capabilities of AI—such as continuous monitoring, anomaly detection, and uncertainty management—are combined with the human's cognitive flexibility, including abductive reasoning, intuition, and context-sensitive decision-making. Resilience in ATC systems implies distributed cognitive capability.

The human retains the ability for flexible adaptation and situated learning, while AI contributes computational speed, amplified memory, and continuous system surveillance. Together, these elements constitute a system that not only enables reaction to unforeseen events, but also allows for anticipation and continuous reconfiguration of operational strategies. Such a model demands careful interface design, training programs centered on human–machine collaboration, and ethical governance of intelligent systems. Capra [4] emphasizes that life itself is a systemic process characterized by interconnections and continuous flows of information; analogously, resilient human–AI teaming can be conceptualized as a living network in which safety is not guaranteed by individual components, but by the quality of interactions, mutual adaptability, and shared situation awareness [1]. This paradigm acknowledges that safety is a continuous process of co-evolution among technology, environment, and human agents—requiring distributed capacities for perception, interpretation, and action.

Air Traffic Control thus represents an ideal testbed for studying the transformations of cognition in the age of AI. The high-stakes nature of this environment, coupled with the necessity for coordination among multiple agents and the growing complexity of information flows, renders it an exemplary field for investigating both the possibilities and constraints of human–AI collaboration.

Acknowledgments

The authors wish to thank the anonymous reviewers for their insightful feedback, which greatly improved the quality of this paper. This work is part of the SESAR CODA project; project type: exploratory research; SESAR program: Digital European Sky; grant ID: 101114765.

Declaration on Generative AI

During the preparation of this work, the author used GPT-4 to check grammar and spelling and to assist in summarizing parts of the content. After using these tools, the author carefully reviewed and edited the material as needed and takes full responsibility for the final content of the publication.

References

[1] E. Hollnagel, D. D. Woods, N. Leveson (Eds.), Resilience Engineering: Concepts and Precepts, Ashgate, Aldershot, 2006.

[2] S. Dekker, S. Pruchnicki, Drifting into failure: theorising the dynamics of disaster incubation, Theoretical Issues in Ergonomics Science 15 (2014) 534–544.

[3] L. Save, Un colpevole ci dovrà pur essere. I luoghi comuni sugli incidenti e le strategie più efficaci per evitarli, Primiceri, Padova, 2019.

[4] F. Capra, The Web of Life: A New Scientific Understanding of Living Systems, Anchor Books, New York, 1996.

[5] D. H. Meadows, Thinking in Systems: A Primer, Chelsea Green Publishing, White River Junction, 2008.

[6] D. Kahneman, Thinking, Fast and Slow, Farrar, Straus and Giroux, New York, 2011.

[7] S. Barrena, J. Nubiola, Abduction: The logic of creativity, in: Bloomsbury Companion to Contemporary Peircean Semiotics, Bloomsbury, 2019, pp. 185–203.

[8] D. A. Schön, The Reflective Practitioner: How Professionals Think in Action, Basic Books, London, 1991.

[9] A. M. Madni, S. Jackson, Towards a conceptual framework for resilience engineering, IEEE Systems Journal 3 (2009) 181–191.

[10] C. Janiesch, P. Zschech, K. Heinrich, Machine learning and deep learning, arXiv preprint arXiv:2104.05314 (2021).

[11] J. B. Morrison, R. L. Wears, Modeling rasmussen's dynamic modeling problem: drift towards a boundary of safety, Cognition, Technology & Work 24 (2022) 127–145.

[12] T. Miller, Explanation in artificial intelligence: Insights from the social sciences, Artificial Intelligence 267 (2019) 1–38.

[13] J. D. Lee, K. A. See, Trust in automation: Designing for appropriate reliance, Human Factors 46 (2004) 50–80.

[14] B. Shneiderman, Human-Centered AI, Oxford University Press, 2022.

[15] E. Hollnagel, D. D. Woods, Cognitive systems engineering: New wine in new bottles, International Journal of Man-Machine Studies 18 (1983) 583–600.

[16] G. Klein, D. D. Woods, J. M. Bradshaw, R. R. Hoffman, P. J. Feltovich, Ten challenges for making automation a "team player" in joint human-agent activity, IEEE Intelligent Systems 19 (2005) 91–95.

[17] E. Cappuccio, A. Esposito, F. Greco, G. Desolda, R. Lanzilotti, Explanation user interfaces: A systematic literature review, arXiv preprint arXiv:2505.20085 (2025).

[18] M. R. Endsley, Toward a theory of situation awareness in dynamic systems, Human Factors 37 (1995) 32–64.

[19] R. Parasuraman, V. Riley, Humans and automation: Use, misuse, disuse, abuse, Human Factors 39 (1997) 230–253.

[20] E. Hutchins, Cognition in the Wild, MIT Press, Cambridge, MA, 1995.

[21] D. Wang, Q. Yang, A. Abdul, B. Y. Lim, Designing theory-driven user-centric explainable ai, in: Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, ACM, New York, 2020, pp. 1–15.

[22] SESAR Joint Undertaking, European atm master plan, https://www.sesarju.eu/masterplan, 2025. Retrieved August 2025.

[23] CODA Project, Coda progress and outcomes, https://iptc.upm.es/coda/results/, 2025. Retrieved August 2025.