Cognition, Technology & Work
https://doi.org/10.1007/s10111-026-00884-3

RESEARCH

Eliciting operational requirements for transparent adaptive automation strategies in air traffic control

Alexandre Veyrié[1] · **Sara Kebir**[1] · **Guillaume Truong**[1] · **Laurie Marsman**[2] · **Christophe Hurter**[1]

Received: 11 March 2026 / Accepted: 8 June 2026
© The Author(s) 2026

Abstract

Adaptive automation is increasingly explored in air traffic control as a means to manage variability in operational demand while preserving controller engagement, situation awareness, and authority. This paper reports the COntroller adaptive Digital Assistant (CODA) contribution on making adaptive task delegation governable in supervisory work. We specify an adaptive automation strategy for non-critical air traffic control support tasks that combines explicit mode logic, authority contracts, and interface mechanisms that render delegation inspectable at the point of action. Adaptation is formalised as a closed-loop policy driven by three trigger classes—traffic demand, task/load demand, and an operator-readiness dimension—computed over both current and short-horizon predicted states to anticipate demand peaks and to support timely, bounded transitions. The strategy operationalises three graduated proactivity modes with explicit authority management: Manual (monitor-and-inform), *Proactive-Light* (recommend-and-confirm), and *Proactive-Strong* (initiate-and-notify with veto), designed to support appropriate reliance and maintain continuous legibility of task ownership. We detail delegation rules that constrain automation to supportive functions, and we present a human–machine interface that externalises automation state, proposed actions, and ownership while incorporating concise, actionable rationale cues intended to support controllability with minimal overhead. Finally, structured workshops with expert controllers are synthesised to derive operational acceptability constraints on delegation boundaries, transition timing, preview horizons, and explanation selectivity, grounding the design in the realities of air traffic control work.

Keywords Adaptive digital assistant · Explainable AI · Human-AI teaming · Air traffic management

1 Introduction

Artificial Intelligence (AI) is increasingly regarded as a means to enhance the performance of Air Traffic Control Officers (ATCOs) by supporting information processing, coordination, and routine yet workload-critical activities (Islam et al. 2022; Degas et al. 2022; Cocchioni et al. 2023; Malakis et al. 2023; Ortner et al. 2022). In high-tempo, high-reliability socio-technical domains such as Air Traffic Management (ATM) and Air Traffic Control (ATC), however, the central challenge extends beyond the mere algorithmic feasibility of AI solutions. The core issue concerns the governability of automation in context: whether task allocation, authority distribution, and control transitions remain transparent, predictable, and recoverable at the moment of interaction, without imposing additional monitoring burdens on human operators (Feigh et al. 2012; Kaber and Endsley 2004; Wickens 1999).

✉ Alexandre Veyrié
alexandre.veyrie@enac.fr

✉ Christophe Hurter
christophe.hurter@enac.fr

Sara Kebir
fs.kebir@gmail.com

Guillaume Truong
guillaume.truong@gmail.com

Laurie Marsman
laurie.marsman@nlr.nl

[1] Fédération ENAC ISAE-SUPAERO ONERA, Université de Toulouse, 31400 Toulouse, France

[2] Department of Safety and Human Performance, Royal NLR - Netherlands Aerospace Centre, Amsterdam
1059, The Netherlands

Published online: 23 June 2026

This challenge is exacerbated by well-documented human–automation failure modes. Inadequately designed or poorly governed automation can induce operator complacency, diminished situational awareness, skill degradation, and out-of-the-loop performance problems, especially when task allocation mechanisms are opaque, authority boundaries are ill-defined, or control transitions are insufficiently supported (Endsley and Kiris 1995; Endsley 2017; Wickens 1999; Berberian et al. 2017; Svensson et al. 2021; Sebastiani et al. 2020; Lee and See 2004). Consequently, the principal scientific and design objective is not simply to increase the degree of automation, but to render AI proactivity controllable. Adaptive delegation policies must be auditable and explainable, while ensuring that human operators retain meaningful authority over safety-critical decisions.

Rising traffic demand and complexity, combined with capacity limitations and staffing constraints (EUROCONTROL 2024, 2022), are accelerating the transition toward higher levels of automation in ATC. Within the SESAR (Single European Sky ATM Research) roadmap (SESAR Joint Undertaking 2025; Bolić and Ravenhill 2021), AI-enabled automation is being actively investigated to support ATCOs across a broad range of operational tasks and to facilitate effective Human–AI Teaming (HAIT) (Gerdes et al. 2022; Jameel et al. 2023; Hunger et al. 2024; Malakis et al. 2023).

The COntroller adaptive Digital Assistant (CODA) project, funded by the SESAR 3 Joint Undertaking, addresses this challenge by developing a human-centred Adaptive Automation (AA) concept for non-critical ATC support tasks. In this paper, HAIT is not treated as an achieved property of CODA, but as a design horizon. It helps specify the interaction conditions under which adaptive delegation may become operationally acceptable: explicit task ownership, observable automation state, predictable transitions, human approval or veto, and concise explanations of system behaviour. Accordingly, the manuscript does not aim to demonstrate mature HAIT performance. Rather, it reports how HAIT-relevant requirements were translated into an adaptive delegation policy and Human–Machine Interface (HMI) mechanisms through a user-centred design process.

In the remainder of the paper, CODA denotes both the project and, when clear from context, the resulting digital assistant concept and prototype implementing the adaptive delegation logic and associated HMI mechanisms. CODA is grounded in the AA literature, particularly levels of automation and dynamic function allocation, to distribute non-critical tasks between the operator and the assistant in a context-dependent way (Parasuraman 2000; Kaber and Endsley 2004; Bernabei and Costantino 2024; Nylin 2023; Rooij et al. 2024). The objective is not to displace the ATCO, but to preserve human authority over safety-critical separation while reducing avoidable cognitive overhead and stabilising workload during demand peaks (Metzger and Parasuraman 2005; Bernabei and Costantino 2024). From this perspective, CODA's core scientific problem is calibration: when assistance should be offered, how proactive it should be, and how the interface should communicate capability, uncertainty, rationale, and automation state so that ATCOs maintain accurate mental models and appropriate reliance (Amershi et al. 2019; Kirwan 2025; Hoffman et al. 2018; Balfe et al. 2015; Gerdes et al. 2025).

A distinctive feature of CODA is that its adaptive policy considers both operational demand, through traffic and task-load indicators, and operator-state indicators as candidate inputs for adaptation. Contemporary neuroergonomics research suggests that physiological signals such as Electroencephalography (EEG), Electrocardiography (ECG), Electrooculography (EOG), and Electrodermal Activity (EDA) can support real-time inference of workload, fatigue, stress, and vigilance, which are directly relevant to out-of-the-loop risk and automation handover timing (Berberian et al. 2017; Dehais et al. 2020; Hurter et al. 2026; Li et al. 2021). Passive brain-computer interfaces and neurophysiologically informed AA have therefore been proposed to adapt assistance according to inferred operator readiness, with the aim of maintaining engagement and reducing out-of-the-loop degradation as proactivity increases (Flumeri et al. 2019; Aricò et al. 2016, 2020). CODA builds on these foundations to operationalise a neurophysiologically grounded, human-centred assistant whose adaptations are intended to remain understandable, predictable, and manageable by the ATCO (Amershi et al. 2019; Causse et al. 2025; Endsley 2017; Langan-Fox et al. 2009).

For this reason, CODA treated end-user and stakeholder workshops as a core input to system specification, not as a late-stage validation add-on. In high-stakes settings, user-centred design is essential to identify which functions are acceptable for delegation, define automation boundaries, shape explanations and feedback that preserve situation awareness, and reveal organisational constraints related to training, accountability, and safety culture (Amershi et al. 2019; Bernabei and Costantino 2024). The present study therefore consolidates the theoretical foundations and practical implementation of CODA's AA framework for en-route ATC. It reports how a mode-based delegation policy was translated into interface requirements, interaction mechanisms, and representative mock-ups, up to a high-fidelity prototype and final HMI. Complementary work within the project addresses signal processing, model training, and neurophysiological validation of human-state inputs, including psychophysiological and behavioural metrics, as triggering dimensions for automation (Ronca et al. 2024; Giorgi et al. 2025; Dello Iacono et al. 2025; Ricci et al. 2025).

Importantly, this paper focuses on the user-centred design and operational specification of CODA's adaptive HMI, rather than on the quantitative validation of CODA's performance effects. Thus, the present manuscript addresses how adaptive task delegation was translated into HMI requirements, how controller feedback shaped the interface, and how workshop-derived themes constrained authority, transparency, and task-ownership mechanisms. The expected benefit of CODA should therefore be interpreted as a design-level and proximal human-factors benefit, not as a demonstrated system-level capacity gain. By supporting the delegation of bounded, non-critical, workflow-relevant tasks, CODA is intended to reduce routine task-management overhead and preserve controller resources for human-led activities, including traffic monitoring, coordination, tactical planning, separation assurance, and conflict resolution. A complementary CODA paper reports the human–machine teaming evaluation and preliminary human-in-the-loop validation results (Bonelli et al. 2025). The present paper addresses the complementary design question: how the AA concept was made operationally legible, controllable, and acceptable through a user-centred HMI design process.

To guide the scope and interpretation of the paper, we address three design-oriented research questions. They are framed around three criteria central to AA in ATC: *controllability*, referring to the preservation of ATCO authority and override capability; *transparency*, referring to the visibility of automation state, task ownership, and system rationale; and *operational fit*, referring to the compatibility of the delegation logic with controller workflow, task criticality, and workload constraints. **RQ1** – How can an AA policy be specified so that proactive delegation of non-critical ATC support tasks remains operationally bounded, predictable, and controllable by the ATCO? **RQ2** – Which HMI mechanisms can make automation state, task ownership, transition behaviour, and system rationale sufficiently explicit to support calibrated reliance with minimal additional cognitive overhead? **RQ3** – What constraints and expectations do ATCOs articulate regarding delegation boundaries, authority management, and automation-level transitions in a supervisory control setting?

In the following, **RQ1** is addressed through the conceptual framing and policy formalisation of CODA's AA strategy (Sects. 2, 3 and 4); **RQ2** through the HMI mechanisms supporting transparency, approval, veto, and task ownership (Sect. 5); and **RQ3** through the workshop-derived constraints and acceptance conditions elicited from expert ATCOs (Sect. 6).

Accordingly, this paper makes three contributions. **C1** – *Mode-logic and delegation specification*: we specify CODA's adaptive delegation policy for non-critical ATC support tasks, including the rationale for three graduated proactivity modes and the associated delegation rules, as an implementable policy artefact. **C2** – *Interface-mediated authority and transparency model*: we operationalise authority management through HMI mechanisms, including approval, veto, explicit task ownership, and transition feedback, designed to preserve ATCO primacy and support calibrated reliance. **C3** – *Workshop-grounded requirements and constraints*: we report qualitative insights from structured engagements with expert ATCOs, synthesised into design requirements and operational constraints that delimit delegation boundaries and interaction behaviour in a supervisory control setting.

2 Adaptive automation and human-AI teaming

In dynamic automation, the allocation of functions may change over time. When the system initiates the allocation change, the mechanism is usually described as adaptive automation; when the human initiates it, it is described as adaptable automation; and when allocation authority can be exercised by either agent under defined constraints, the configuration is hybrid (Bernabei and Costantino 2024; Inagaki 2003; Kaber and Endsley 2004; Scerbo 2018). Adaptive Automation thus denotes the real-time modulation of automation level and/or function allocation according to the operational context and the human operator's state, with the aim of maintaining safety and performance across varying task demands (Bernabei and Costantino 2024; Parasuraman et al. 2000; Kaber and Endsley 2004). Foundational research frames this issue through the "automation conundrum". As automation becomes more capable and autonomous, the human role may shift toward passive supervision, increasing the risk of reduced situation awareness, degraded vigilance, loss of manual readiness, and out-of-the-loop performance problems (Berberian et al. 2017; Endsley and Kiris 1995; Endsley 2017).

Adaptive automation should therefore not be understood as simply increasing automation under high workload. Rather, it is a closed-loop socio-technical control problem requiring explicit specification of which functions adapt, such as information acquisition, analysis, decision support, or action execution; when and why adaptation occurs, including triggers and thresholds; and how changes in authority and responsibility are made transparent to the operator (Endsley 2023). Within CODA's ATC framework, this implies that adaptive policies must rely on observable and measurable indicators, such as workload, attention, fatigue, and task tempo, while preserving the ATCO's ability to supervise, anticipate, and intervene, rather than merely increasing system throughput.

Human–AI teaming extends this perspective by conceptualising AI-based automation not solely as an instrumental tool, but as a teammate-like agent engaged in coordinated joint activity with human operators. Human–AI teaming can be defined as a dynamic, interdependent mode of cooperation in which one or more humans and one or more AI-based systems coordinate their actions, exchange task-relevant information, and adapt roles or strategies in pursuit of shared objectives (Chen et al. 2025; Human-AI Teaming 2022; Pham et al. 2024). Within this framework, overall team performance cannot be reduced to the technical accuracy of the AI component. Instead, it is an emergent property of team cognition and interaction processes, including communication, coordination, backup behaviour, role clarity, and adaptive task allocation.

Effective HAIT consequently depends on clearly specified authority boundaries, AI behaviour that is both predictable and steerable, appropriately calibrated trust that mitigates both misuse and disuse, and shared situation awareness supported by transparency and explainability mechanisms (Amershi et al. 2019; Chen et al. 2025; Endsley 2023; Hancock 2022). The construct of shared mental models is particularly salient in this context, as effective teams rely on aligned expectations regarding task goals, teammates' capabilities and limitations, and their likely behaviours (Andrews et al. 2023). In human–AI teams, however, achieving such alignment is complicated by learning-based system architectures, non-stationary performance profiles, and interaction that is predominantly mediated through user interfaces.

In aviation, practitioners therefore emphasise that HAIT remains, in part, an aspirational construct: many currently deployed or emerging systems still exhibit"design debt" in the form of automation overreliance, opacity, and insufficiently supported coordination. Realising robust HAIT thus requires explicit operational concepts, well-defined interaction protocols, and systematic empirical validation, rather than reliance on algorithmic performance metrics alone (Luciani et al. 2020; Van Droogenbroeck et al. 2025).

The requirements for effective AA and HAIT converge around three design and evaluation imperatives. First, adaptation triggers must be diagnostic, reliable, and stable. They should ideally combine task- and context-related cues with operator-state indicators, including physiological measures, to avoid oscillatory or unexpected adaptations that may undermine trust and accountability (Aricò et al. 2017; Flumeri et al. 2019; Bernabei and Costantino 2024; Chancey et al. 2025; Sebastiani et al. 2020). Second, teaming requires transparency mechanisms that clarify what the AI is doing, what it perceives, and what it is likely to do next, together with explainability mechanisms that support mental model formation by conveying why the AI recommends or executes a given action. In high-tempo operations, the timing and information bandwidth of these explanations must be carefully managed to avoid unnecessary cognitive load (Xie et al. 2021; Chen et al. 2024, 2025; Cartocci et al. 2026). Third, teaming quality cannot be inferred from task performance alone. Validation should therefore combine task outcomes, such as accuracy, time, and errors; human factors indicators, such as workload, situation awareness, fatigue, vigilance, and stress; and team-level interaction measures, such as communication effectiveness, role allocation, adaptability, and trust dynamics (Damacharla et al. 2018; Lashley et al. 2019). In CODA, these considerations motivate a design stance in which AA is treated not only as an internal control function, but as a teamwork enabler. Adaptations must be intelligible to the human, mediated through the interface, and assessed according to their impact on joint human-AI performance rather than isolated automation gains (Kirwan 2024).

Recent aviation guidance provides an additional regulatory and human-factors frame for the type of HAIT addressed in CODA. The EASA AI Concept Paper Issue 2 distinguishes Level 2 AI applications as systems involving human-AI cooperation or collaboration under human oversight, with attention to shared operational authority, operational explainability, end-user trust, and the definition of task allocation between the human operator and the AI-based system (European Union Aviation Safety Agency 2024). This framing is directly relevant to AA in ATC, because the distinction between predefined and dynamic allocation patterns determines not only what the automation may do, but also how authority, oversight, intervention, and responsibility remain intelligible to the controller. In particular, Level 2B human-AI collaboration is associated with the capacity to readjust strategies and task allocation in real time while maintaining communication and shared situation representation between the human and the AI-based system (European Union Aviation Safety Agency 2024).

Within this perspective, CODA should not be interpreted as a mature Level 2B operational system, but as an early design- and requirements-oriented exploration of transparent adaptive task delegation for non-critical ATC support tasks. Its mode logic deliberately constrains the delegation envelope, preserves ATCO primacy for safety-critical decisions, and makes task ownership, mode transitions, and intervention possibilities explicit through the HMI. This conservative interpretation is consistent with EASA's emphasis on documenting task allocation patterns, preserving human oversight, and ensuring that end users receive understandable and timely information about AI-based system behaviour.

3 COntroller adaptive Digital Assistant (CODA)

The COntroller Adaptive Digital Assistant (CODA) is conceived as a human-centred AA concept intended to support en-route ATCOs in the management of bounded, non-critical, workflow-relevant tasks, while preserving the ATCO's exclusive responsibility for safety-critical control, notably separation assurance and conflict resolution. CODA operationalises adaptive (and, where appropriate, adaptable) automation through a closed-loop cycle that combines (i) continuous sensing of operational context and operator state, (ii) predictive inference of upcoming task demand and cognitive readiness, and (iii) transparent, mode-based assistance that reallocates selected non-critical tasks under explicit constraints on authority, safety, stability, and recoverability.

Several terms are used in a specific sense throughout the paper. The *adaptive policy* refers to the rules by which CODA selects an automation mode from operational context, task-load indicators, and operator-state estimates. *Mode logic* refers to the transition rules governing movement between *Manual*, *Proactive-Light*, and *Proactive-Strong* modes. *Delegation logic* refers to the rules determining which non-critical tasks may be assigned to the ATCO or to CODA, and under which approval, confirmation, or veto conditions. Finally, *proactivity* denotes the system's ability to anticipate near-future demand and initiate a recommendation or bounded action before performance degradation occurs, while preserving the controller's authority to approve, reject, or override automation.

3.1 Conceptual scope and safety boundaries

CODA is designed around a dual-function view of ATC work: (1) collision prevention (critical, non-delegable), and (2) orderly flow management (adaptable, partially delegable). Accordingly, CODA constrains automation to tasks that can improve efficiency and workload regulation without directly influencing separation assurance. This scope limitation is key to mitigating authority migration, inappropriate reliance, and the erosion of monitoring competence associated with high automation. In CODA, the ATCO therefore remains the final decision authority, and CODA assistance is designed to remain interruptible and reversible such that manual reversion is feasible without inducing unsafe transitions (Endsley 2017; Hancock 2022).

3.2 System architecture and information flow

CODA is implemented as a modular architecture coupling operational data streams, human-state sensing, and adaptive decision logic, consistent with models that distinguish automation across functional stages (information acquisition, information analysis, decision selection, and action implementation) Parasuraman et al. (2000); Parasuraman (2000). At runtime, the system is organized into four functional layers:

(1) **Operational context acquisition**: CODA does not ingest undifferentiated traffic data, but structured descriptors of the current and near-future operational situation. These include traffic-demand descriptors (e.g., number of aircraft under responsibility, sector entry/exit events, convergence or sequencing situations, and predicted traffic evolution), task-event descriptors derived from the ATCO task taxonomy (e.g., radar contact, flight release, SSR code change, STAR allocation, provision of standard information, coordination or handover events), and interaction traces produced through the controller working position and CODA HMI. In the present paper, the term *ATC interactions* refers primarily to logged controller-position events, task-management actions, and HMI interactions, rather than to unrestricted natural-language analysis of radio communication. These signals are treated as inputs for estimating traffic demand and task/interaction load, and for identifying when non-critical support tasks may become eligible for recommendation or delegation.

(2) **Human-state estimation and prediction module**: CODA uses the output of a human-state sensing pipeline developed in the broader CODA project (Ronca et al. 2024; Giorgi et al. 2025; Dello Iacono et al. 2025; Ricci et al. 2025). This pipeline combines unobtrusive physiological and behavioural indicators to estimate cognitive states relevant to AA, including workload, stress, fatigue, vigilance, and drowsiness. Depending on the experimental configuration, such indicators may include neurophysiological, cardiovascular, electrodermal, ocular/behavioural, voice, and interaction-based features. In the present manuscript, these signals are not modelled directly. They are used only through the derived Mental State/Readiness index M, discretised into Low, Medium, or High, together with a confidence estimate c_M. This distinction is important: the paper does not validate the human-state models themselves, but specifies how their categorical output can be integrated conservatively into adaptive delegation logic and HMI design.

(3) **Adaptive Automation Strategy (AAS) engine**: The AAS integrates current and predicted indicators (traffic/task demand and operator state) to select an assistance policy and an automation mode, i.e., an operationalised "level of automation" suitable for the ongoing context.

CODA's AAS is designed to counteract overload-related degradation while avoiding excessive or unstable mode shifts that can erode mode awareness and predictability. The AAS is further constrained by HAIT requirements: calibrated trust, meaningful human control, and consistent handover conditions.

(4) **Interaction and execution layer**: CODA assistance is delivered via an HMI designed to preserve mode awareness, task-status visibility, and contestability. The interface is designed to communicate current mode, task allocation status, and automation actions in a way that should support rapid inspection in time-pressured operations. Whether this produces faster understanding remains an empirical question for dedicated evaluation that is out of the scope of the current paper. For explainability requirements specific to ATM, CODA aligns with the broader explainable AI literature emphasizing transparency and operator-centered interpretability, while acknowledging the need for evaluation criteria for explanatory effectiveness.

3.3 Task model, delegation targets, and criticality boundary

The delegation logic underlying CODA is grounded in a task taxonomy derived from routine operational events in en-route control (Zamarreño Suárez et al. 2023, 2022; Gutiérrez Teuler et al. 2022; Munoz-De-Escalona et al. 2025). This taxonomy includes events such as radar contact, sector entry/exit within a shared flow, flight release, top-of-climb/top-of-descent notifications, IFR SSR code change, STAR allocation, provision of essential information such as QNH, navigation-related events, procedural clearances, flow-related actions, coordination/handovers, and efficiency-oriented adjustments such as weather-related detours. These events are used both to identify potential support opportunities and to estimate workload, coordination demand, and traffic complexity. However, their inclusion in the taxonomy does not imply that all associated control actions are delegable.

A central constraint of CODA is that it targets non-critical support tasks only. In the present design scope, a task is considered eligible for CODA support only when the following conditions are jointly satisfied: (i) it does not directly create, modify, or replace a separation-related or safety-critical clearance; (ii) it has no direct and immediate impact on tactical separation minima; (iii) it is bounded in time and consequence; (iv) it remains visible, interruptible, reversible, and recoverable by the ATCO; and (v) the local traffic, coordination, and communication context does not make delegation operationally ambiguous. Eligibility is therefore context-dependent rather than permanently attached to a task label.

For example, a speed adjustment used to maintain flow may be eligible for CODA support, whereas a speed adjustment used for separation remains ATCO-owned. Similarly, a direct route used to shorten a flight plan may be considered only when conflict-free, whereas a direct route used to solve or avoid a conflict is outside CODA's delegable scope. Events associated with tactical separation or queuing, such as vectoring, level or speed changes for separation/queue management, holding clearances, conflict detection/resolution, emergency handling, and any instruction whose timing or content is directly linked to maintaining separation minima, remain strictly non-delegable. Illustrative examples of CODA-eligible support tasks and delegation boundaries are presented in Table 1.

Conflict detection is treated particularly conservatively. Although conflict probes and alerts may exist in contemporary ATC systems and may contribute to the controller's traffic picture, CODA does not treat conflict detection as a delegable task in the present paper. The reason is that conflict detection is operationally coupled with separation monitoring, tactical intent, and the selection of conflict-resolution actions. A conflict-related cue may therefore be used as a contextual signal that increases the Traffic/Complexity index or prevents delegation of otherwise routine tasks, but CODA does not assume ownership of conflict detection, separation assurance, or conflict resolution. Flights involved in a detected or suspected conflict are excluded from CODA task takeover and remain under ATCO control.

Within this boundary, CODA structures assistance through a small set of operational modes intended to support stable mental models, reduce mode confusion, and ensure predictable behaviour. The expression *CODA assumes a task* therefore refers only to assistant-side task-management ownership within the prototype HMI: the task is moved to the CODA side of the interface, marked as being handled by the assistant, and associated with a notification, rationale, and veto/reclaim possibility. It does not imply that CODA is authorised to issue separation-related clearances, communicate unrestricted instructions to pilots, replace ATCO tactical judgement, or act outside a human-as-final-authority interaction contract.

4 Adaptive automation strategies

Sections 2 and 3 motivated AA as a HAIT problem in which authority, task ownership, and transitions of control must remain governable at interaction time. We now operationalise these principles in CODA by specifying (i) a small set of proactive modes with explicit authority contracts,

Table 1 Illustrative examples of CODA-eligible support tasks and delegation boundaries

Task example	Current ATCO-centred interaction	CODA support in the prototype	Delegation boundary
SSR code change	The ATCO identifies the need for a transponder-code change, issues or prepares the instruction, monitors acknowledgement, and updates the task state	CODA may detect that the code-change task is pending, display it as eligible, remind the ATCO, or prepare the task for approval. In *Proactive-Light*, execution requires ATCO confirmation; in *Proactive-Strong*, CODA may assume task-management ownership with notification and veto/reclaim	Not eligible if the code change is part of an abnormal, degraded, or safety-critical situation requiring immediate ATCO judgement
STAR assignment or standard arrival preparation	The ATCO assigns or prepares a standard arrival route according to traffic flow, destination, and coordination constraints	CODA may recommend the standard assignment or mark it as a pending flow-management task, making the affected flight and rationale visible	Not eligible when the assignment is directly linked to tactical conflict resolution, unstable sequencing, or local constraints not represented in the system
Standard information provision	The ATCO provides routine information, such as relevant weather, equipment, or transfer-condition information, when required by the operational context	CODA may notify the ATCO that standard information should be provided, prepare a reminder, or display the affected flight/task in the assistant panel	Not eligible when the information is safety-critical, ambiguous, or requires immediate judgement about abnormal conditions
Coordination or handover preparation	The ATCO coordinates with adjacent sectors, receives/retransmits authorisations, and ensures that transfer conditions are understood	CODA may monitor pending coordination/handover events, remind the ATCO, or prepare the task for confirmation	Not eligible when there is unresolved inter-sector negotiation, unusual transfer conditions, or radio/coordination ambiguity
Flow-maintenance speed or route suggestion	The ATCO may use speed or route adjustments to smooth traffic flow or reduce inefficiencies when no separation issue is involved	CODA may recommend a flow-oriented adjustment or identify the task as suitable for assistant support	Not eligible when the adjustment is used for separation, conflict avoidance, emergency handling, or when local tactical intent makes automation inappropriate

The examples are not intended as an exhaustive operational procedure, but as a clarification of how non-critical task support is interpreted in the present HMI-design paper

(ii) a triggering logic grounded in traffic complexity, task demand, and (when available) human-readiness cues over current and predicted horizons, and (iii) interface mechanisms that externalise mode state, task ownership, and bounded rationales at the moment of delegation. This specification constitutes the core contribution addressing **RQ1**, and it provides the basis for the HMI mechanisms (**RQ2**) and workshop-derived constraints (**RQ3**) reported later.

Building on the concepts introduced in Sect. 2 and the CODA scope in Sect. 3, we formulate CODA's Adaptive Automation Strategy (AAS) as a closed-loop delegation policy for non-critical support tasks (CODA Consortium 2025; Marsman et al. 2024). The strategy maps a compact set of operational and human-state indices—traffic complexity, task demand, and human readiness, each considered over current and predicted horizons—to one of three proactive modes and associated delegation actions. The emphasis is not on proposing a new definition of AA, but on specifying an auditable, interaction-time governable policy that preserves ATCO authority through explicit approval, veto, and takeover mechanisms.

4.1 Proactive modes and triggering indices

Adaptive automation is operationalised through three graduated modes with explicit authority management. In *Manual* mode, the ATCO retains full control while the assistant monitors and informs. In *Proactive-Light* mode, the assistant recommends taking over selected non-critical tasks, requiring explicit ATCO approval to support appropriate reliance. In *Proactive-Strong* mode, the assistant initiates the delegation of selected non-critical tasks, while preserving ATCO authority through immediate notification and a short veto/override window, limiting surprise while enabling timely relief under high demand.

The three-mode structure was chosen as a compromise between operational simplicity and graduated proactivity. A larger number of automation levels would allow finer adaptation but would increase the risk of mode confusion and make it harder for ATCOs to maintain a stable mental model of CODA's behaviour. Conversely, a binary manual/automated distinction would not provide enough granularity to distinguish recommendation, confirmation, and bounded automation initiation. The three modes therefore map onto three authority contracts that are simple enough to remain inspectable: *Manual* preserves full ATCO task ownership; *Proactive-Light* introduces assistant initiative while requiring explicit approval; and *Proactive-Strong* allows bounded initiative under notification, veto, and reclaim conditions.

Table 2 Role of the three CODA indices in adaptive delegation

Index	Example input classes	Treatment in the AAS	Role in mode adaptation
T: Traffic & complexity	Aircraft count and density, sector entry/exit events, convergence/sequencing situations, predicted traffic evolution, local traffic complexity indicators	Transformed into current and predicted ordinal levels $T_c, T_p \in \{\text{Low, Medium, High}\}$. A weighted score $S_T = w_c T_c + w_p T_p$ balances current demand and short-horizon anticipation	Indicates operational demand imposed by the traffic situation. High or increasing T can support escalation toward *Proactive-Light* or, when combined with high task demand and reliable readiness information, *Proactive-Strong*
L: Task & interaction load	Logged task events and interaction traces, such as radar contact, flight release, SSR code change, STAR allocation, standard information provision, coordination/handover events, and pending task density	Transformed into current and predicted ordinal levels L_c, L_p. A weighted score $S_L = w_c L_c + w_p L_p$ estimates current and anticipated routine task burden	Indicates whether routine, non-critical task-management demand is increasing. High L favours recommending or initiating delegation of eligible tasks, while preserving ATCO ownership of safety-critical actions
M: Mental state & readiness	Derived output of human-state sensing pipelines estimating workload, stress, fatigue, vigilance, or drowsiness from physiological and behavioural indicators	Transformed into current and predicted ordinal levels M_c, M_p, with an associated confidence estimate c_M. If $c_M < c_{\min}$ or the estimate is unavailable, readiness-driven escalation is disabled	Modulates proactivity only when reliable. High M, interpreted as reduced reserve capacity or increased strain depending on the model convention, can support escalation when T and/or L also indicate demand. It cannot alone authorise unsafe delegation and cannot override task-criticality constraints

The table summarises the information used by the AAS in the present HMI-design paper

Mode transitions are driven by three complementary indices, computed for both *current* ($\cdot_c$) and *predicted* ($\cdot_p$) horizons: a Traffic/Complexity index (*T*), a Task/Interaction Load index (*L*), and a Mental State/Readiness index (*M*). Each is discretised into {Low, Medium, High} to support anticipatory adaptation and interpretable policy design. While T and L are always available from operational context and interaction traces, M is treated as an optional input derived from human-state sensing pipelines (Ronca et al. 2024; Giorgi et al. 2025; Dello Iacono et al. 2025; Ricci et al. 2025); when M is missing or uncertain, CODA degrades to conservative behaviour as specified below. For clarity, M is interpreted here as a support-need indicator derived from mental-state estimates, rather than as a direct clinical or fitness-for-duty label: higher values indicate reduced reserve capacity or increased need for support. More details on the three CODA indices in adaptive delegation are presented in Table 2.

4.2 Three complementary strategies

The three strategies do not represent three separate CODA systems. They intervene at different points of the same adaptive-design process. Strategy 1 is the runtime policy: it specifies how current and predicted indices are transformed into mode transitions during interaction. Strategy 2 is an engineering and audit mechanism: it compresses the 729 possible index configurations into explainable risk envelopes and supports inspection of the policy before deployment. Strategy 3 is a human-grounded calibration mechanism: it uses stakeholder judgements to test whether the proposed mappings between states and adaptation actions are acceptable and understandable. These complementary strategies were introduced because adaptive triggering in ATC must satisfy three simultaneous constraints: it must be computable at runtime, auditable by designers and safety analysts, and acceptable to controllers.

Strategy 1: Threshold-based mode logic (implemented policy). The threshold-based strategy provides an auditable control law that maps the monitored indices to mode transitions while preserving mode awareness through stability safeguards (Fig. 1). Each index $X \in \{T, L, M\}$ is discretised as $\{1, 2, 3\}$ for {Low, Medium, High}, and evaluated over a current (X_c) and predicted (X_p) horizon. For each index we compute a weighted score:

$$S_X = w_c\, X_c + w_p\, X_p, \tag{1}$$

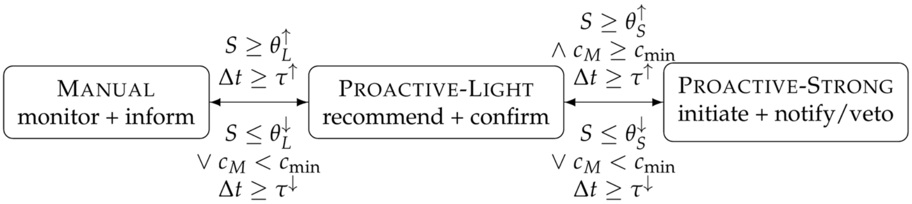

Fail-safe: if the mental-state estimate is missing or low-confidence ($c_M < c_{\min}$), readiness-driven escalation is disabled and the policy degrades to traffic/task triggers only, applying conservative transitions.

Fig. 1 Threshold-based adaptive mode logic in CODA for the delegation of non-critical ATC support tasks. The ATCO-assistant team operates across three graduated modes: *Manual* (*monitor + inform*), *Proactive-Light* (*recommend + confirm*), and *Proactive-Strong* (*initiate + notify/veto*), which differ in the degree of assistant initiative while preserving controller authority. Transitions are driven by a composite support-need score S, with escalation and de-escalation determined by mode-specific thresholds and dwell-time constraints to prevent unstable switching. Escalation to *Proactive-Strong* further requires a sufficiently reliable mental-state estimate ($c_M \geq c_{\min}$); otherwise, the policy adopts a conservative fail-safe and disables readiness-driven escalation. S, composite support-need score; θ_L, threshold for transitions involving *Proactive-Light*; θ_S, threshold for transitions involving *Proactive-Strong*; $\uparrow$ / $\downarrow$, escalation/de-escalation thresholds; Δ t, time during which the condition remains true; τ, minimum dwell time before transition; c_M, confidence in the mental-state estimate; $c_{\min}$, minimum confidence required to authorise readiness-based escalation

with $w_c, w_p \geq 0$ and $w_c + w_p = 1$ to balance responsiveness and anticipation.

Mental-state input is treated as optional. When a mental-state estimate is unavailable or assessed as low-confidence (e.g., $c_M < c_{\min}$), CODA degrades to traffic/task triggers only, and readiness-driven escalation is disabled. In this configuration, proactivity is conservatively capped to *Proactive-Light* (recommend-and-confirm) or maintained in *Manual*, depending on the configured safety policy. No escalation to *Proactive-Strong* is permitted when $c_M < c_{\min}$.

Mode transitions are driven by thresholds with (i) hysteresis (distinct escalation vs. de-escalation thresholds) and (ii) minimum dwell time $\tau_{\min}$ to prevent oscillations. Using an aggregated demand indicator $S(t) = \max(S_T, S_L, S_M)$ (or $S(t) = \max(S_T, S_L)$ when *M* is unavailable/low-confidence), the following applies:

- **Manual** $\rightarrow$ **Proactive-Light** if $S(t) \geq \theta_L^{\uparrow}$ for at least $\tau^{\uparrow}$.
- **Proactive-Light** $\rightarrow$ **Proactive-Strong** if $S(t) \geq \theta_S^{\uparrow}$ for at least $\tau^{\uparrow}$ *and M* is available with $c_M \geq c_{\min}$.
- **Proactive-Strong** $\rightarrow$ **Proactive-Light** if $S(t) \leq \theta_S^{\downarrow}$ for at least $\tau^{\downarrow}$ or if *M* becomes unavailable/low-confidence.
- **Proactive-Light** $\rightarrow$ **Manual** if $S(t) \leq \theta_L^{\downarrow}$ for at least $\tau^{\downarrow}$ or if *M* becomes unavailable/low-confidence.

Typically $\theta_S^{\uparrow} > \theta_L^{\uparrow}$ and $\theta_S^{\downarrow} > \theta_L^{\downarrow}$, with $\tau^{\downarrow} \geq \tau^{\uparrow}$ to favour stable de-escalation.

A key benefit of this design is auditability: each adaptation can be logged and rendered explainable by indicating which index (traffic, task load, and when available mental state) crossed which threshold, thereby supporting calibrated reliance and post-hoc safety assessment.

Strategy 2: Risk-envelope policy implemented as an explainable decision tree. A key challenge for AA in ATC is that the joint state defined by traffic, task demand, and human readiness – each evaluated in both current and predicted horizons – yields a large and potentially opaque decision space. With indices discretised into {Low, Medium, High} for *T*, *L*, and *M* and computed over current and predicted horizons, the resulting state space contains $3^6 = 729$ distinct configurations. While such granularity is useful for capturing realistic combinations of operational and human factors, it is not directly manageable to training, explanation, or operational review.

To make this space operationally tractable and explainable, CODA compresses it into a small number of risk envelopes and associated adaptation actions using an interpretable rule-based model. Concretely, we instantiate this compression through a decision tree classifier trained on the six features $\{T_c, T_p, L_c, L_p, M_c, M_p\}$ and a categorical target action (e.g., *Increase*, *Decrease*, *No Change* – Automation). The tree is deliberately constrained to limited depth (e.g., $\text{max\_depth} = 4$) to avoid over-fragmentation and to preserve human interpretability. Each root-to-leaf path defines a short, auditable "if–then" rule that maps a state configuration to one recommended adaptation. This yields three practical benefits.

First, the tree acts as a risk-envelope compressor. Although the full space contains 729 configurations, the depth-constrained tree partitions it into a small number of leaf regions, each corresponding to a coherent envelope such as low-risk/stable (typically "no change"), moderate-risk/anticipatory (often "increase" to *Proactive-Light* with approval), or high-risk/relief-needed (eligible for escalation toward *Proactive-Strong* when readiness evidence is available and reliable). In other words, the decision tree operationalises the notion of "Low/Moderate/High" risk by learning which combinations of current and predicted indices consistently warrant different degrees of delegation.

Second, the tree provides built-in explanations at the same granularity as the triggering logic. For any adaptation, the system can surface a concise rationale derived from the active path (e.g., "increase automation because traffic complexity is currently high or predicted to be high"), which directly supports transparency, calibrated trust, and shared mental models. This explanation is not a post-hoc add-on: it is structurally tied to the decision rule that produced the action (Fig. 2).

Third, the tree supports engineering review and training. Because the learned policy is expressible both as a visual tree (for inspection) and as text-based rules (for documentation), it can be iteratively refined with stakeholders: designers can adjust feature thresholds, constrain unsafe branches, and reconcile disagreements between expected and learned behaviour without requiring opaque parameter tuning. This is particularly important in CODA's context, where automation must remain controllable and aligned with ATCO authority and where mental-state inputs are optional; branches that rely on *M* can be systematically gated by confidence (Strategy 1), ensuring conservative behaviour when readiness estimates are missing or uncertain.

Overall, Strategy 2 turns a high-dimensional AA problem into a compact set of explainable risk envelopes and actionable rules that are directly compatible with HMI-mediated delegation: it enables adaptation decisions that are (i) operationally tractable, (ii) auditable, and (iii) explainable to ATCOs at the moment of use.

Strategy 3: Space-sampling policy (human-grounded mapping of states to actions). To complement analytic policies with human-grounded preferences, the space-sampling strategy (Fig. 3) builds an empirical mapping from representative (*T*, *L*, *M*) configurations to recommended actions (*Increase*, *Decrease*, *No change*) through structured elicitation. In this exercise, short online questionnaires were completed by CODA consortium partners involved in technical, operational, and human-factors activities, as well as by engineering and controller students and expert ATCOs. The resulting dataset comprised $n = 215$ labelled state-action samples, where each sample corresponds to one judged combination of traffic demand, task/load demand, and controller mental-state/readiness indicators mapped to an adaptation action. These samples were used as preliminary design inputs for inspecting and tuning the adaptive policy,

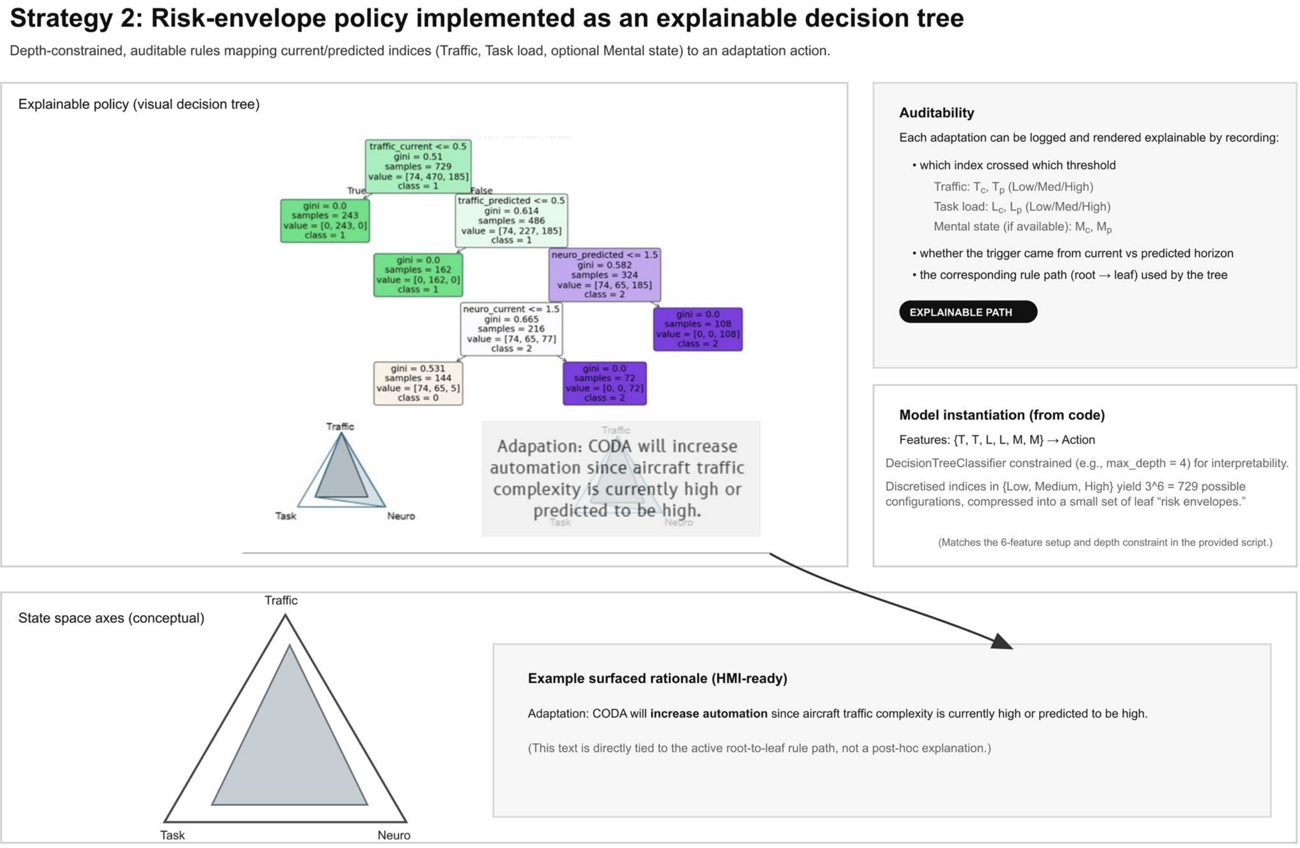
Strategy 2: Risk-envelope policy implemented as an explainable decision tree

Depth-constrained, auditable rules mapping current/predicted indices (Traffic, Task load, optional Mental state) to an adaptation action.

Adapation: CODA will increase automation since aircraft traffic complexity is currently high or predicted to be high.

Auditability

Each adaptation can be logged and rendered explainable by recording:

- which index crossed which threshold
  - Traffic: T_c, T_p (Low/Med/High)
  - Task load: L_c, L_p (Low/Med/High)
  - Mental state (if available): M_c, M_p
- whether the trigger came from current vs predicted horizon
- the corresponding rule path (root → leaf) used by the tree

Model instantiation (from code)

Features: {T, T, L, L, M, M} → Action

DecisionTreeClassifier constrained (e.g., max_depth = 4) for interpretability.

Discretised indices in {Low, Medium, High} yield 3^6 = 729 possible configurations, compressed into a small set of leaf "risk envelopes."

(Matches the 6-feature setup and depth constraint in the provided script.)

Example surfaced rationale (HMI-ready)

Adaptation: CODA will **increase automation** since aircraft traffic complexity is currently high or predicted to be high.

(This text is directly tied to the active root-to-leaf rule path, not a post-hoc explanation.)

Fig. 2 Depth-constrained decision tree implementing CODA's risk-envelope AA policy by mapping the current and predicted indices $\{T_c, T_p, L_c, L_p, M_c, M_p\}$ (each discretised into {Low, Medium, High}) to a categorical adaptation action (*Increase*, *Decrease*, *No Change*)

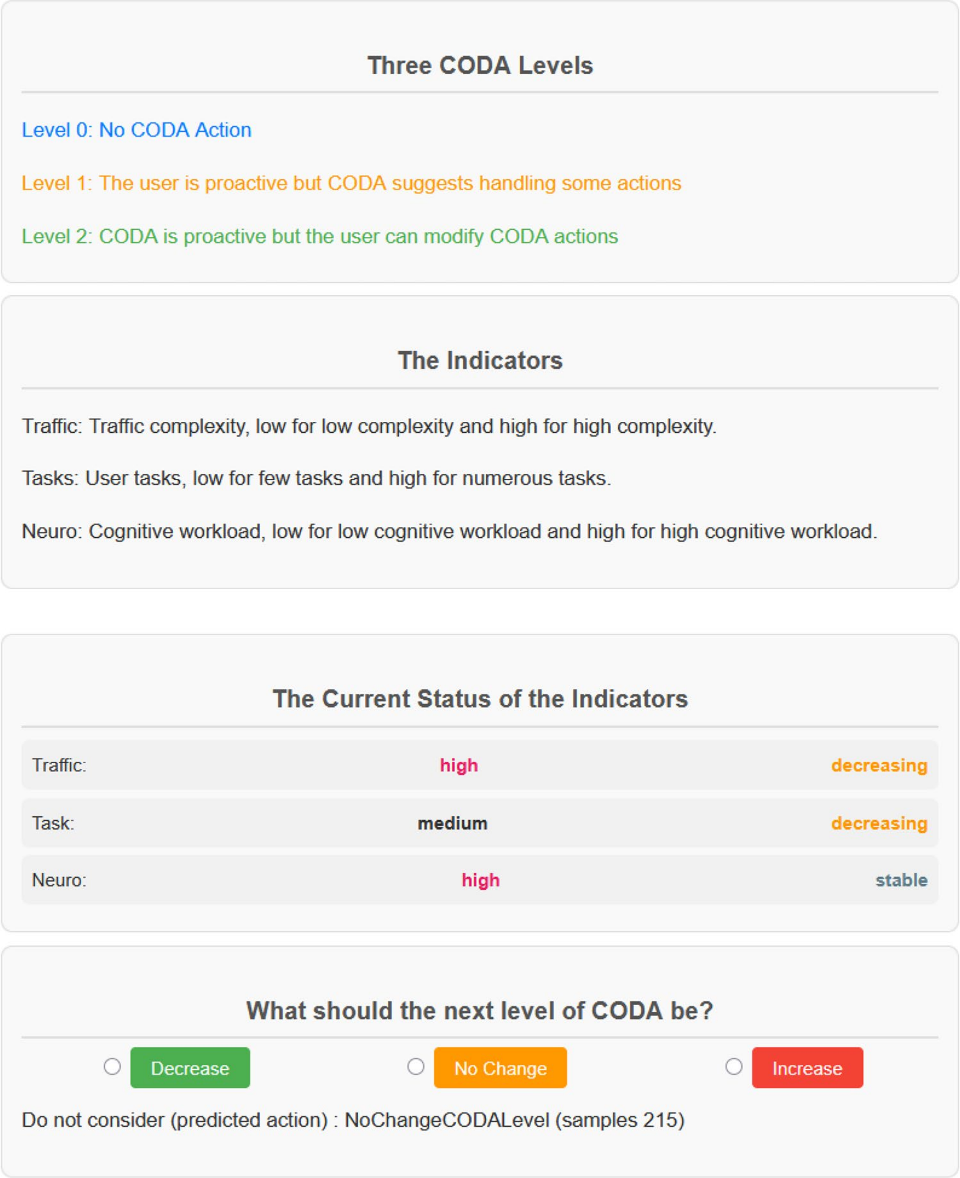

Three CODA Levels

Level 0: No CODA Action

Level 1: The user is proactive but CODA suggests handling some actions

Level 2: CODA is proactive but the user can modify CODA actions

The Indicators

Traffic: Traffic complexity, low for low complexity and high for high complexity.

Tasks: User tasks, low for few tasks and high for numerous tasks.

Neuro: Cognitive workload, low for low cognitive workload and high for high cognitive workload.

The Current Status of the Indicators

Traffic:	**high**	**decreasing**
Task:	**medium**	**decreasing**
Neuro:	**high**	**stable**

What should the next level of CODA be?

○ Decrease ○ No Change ○ Increase

Do not consider (predicted action) : NoChangeCODALevel (samples 215)

Fig. 3 Questionnaires provided to ATM stakeholders displaying the space-sampling policy strategy ($n = 215$ samples) for the AA (Strategy 3). Each sample becomes a labelled example of what stakeholders consider an appropriate adaptation under a given state combination

rather than as evidence of final operational acceptability. They can support the adjustment of thresholds and risk boundaries, reveal mismatches between designer assumptions and stakeholder expectations, and provide traceability for why a given adaptation rule was considered plausible at this stage of development. The primary end-user evidence for operational acceptability and HMI requirements remains the expert ATCO workshops reported in Sect. 6. In this way, space sampling functions as an intermediate bridge between the AA policy layer and the subsequent user-centred HMI design. It informs how adaptations should be presented, justified, and controllable, while leaving their operational validity to be assessed through scenario-based evaluation with controllers.

Illustrative adaptation episode. Consider a short-horizon prediction in which the sector is expected to receive several aircraft entries while two coordination actions and one standard information task are pending. In this case, T_p increases because traffic complexity is predicted to rise, and L_c or L_p increases because several routine task events are active or anticipated. If M is unavailable or low-confidence, CODA can at most recommend support in *Proactive-Light*: the HMI displays the affected task or flight, marks it as eligible for delegation, and presents an approval control for the ATCO. If M is available with sufficient confidence and indicates reduced reserve capacity in combination with high T and L, the policy may escalate to *Proactive-Strong* for eligible non-critical tasks only. The ATCO then sees the active mode, the task ownership change, the affected item, a short rationale such as *predicted traffic and task load high*, and a veto/reclaim option. Separation-related actions remain ATCO-owned throughout the episode.

Implications for the user-centred HMI design. Across all three strategies, allocation changes are treated as interactive transitions rather than silent mode flips: mode status and task ownership must remain visible, automation-initiated changes must be signaled and reversible (veto/override), and explanations should be lightweight and timed to operational tempo (what changed, why, what happens next) (Endsley 2023; Amershi et al. 2019; Chen et al. 2024). These requirements motivate the HMI design choices developed in the next section, in which the interface needs to explicitly support shared goals, transparent delegation, and ATCO agency as primary conditions for safe and effective adaptive teaming (Endsley 2017; Lee and See 2004; Luciani et al. 2020; Pacherie 2008; Sahaï et al. 2019, 2023).

5 User-centred HMI Design

Because the present manuscript reports a user-centred design and requirements process, claims about CODA are formulated as design specifications and intended human-factors effects rather than as demonstrated operational outcomes. Feedback from the workshops was used to assess the plausibility, acceptability, and interpretability of the proposed HMI mechanisms, but not to quantify their effects on response time, workload, situation awareness, or operational performance.

5.1 HMI philosophy and design constraints

CODA's HMI is the primary coordination layer through which adaptive delegation is made observable, controllable, and explainable in a high-tempo supervisory control setting. Its design goal is to support transparent task delegation and controllable automation while aiming to add minimal cognitive overhead to the ATCO's work, notably through lightweight interactions for approval, veto, and status inspection. In line with human-centred and human-AI interaction guidance (Amershi et al. 2019; European Union Aviation Safety Agency 2024), the interface is designed to make the system's state and actions immediately legible (What is the current automation level? Which tasks are owned by whom? What will happen next?), to keep automation behaviour predictable, and to preserve operator authority through straightforward override and veto mechanisms.

A central constraint of CODA is that it focuses on non-critical support tasks and does not perform conflict resolution or other safety-critical control functions. The HMI therefore avoids visual/interaction patterns that could be confused with separation assurance tools, and explicitly differentiates responsibilities and task ownership between the ATCO and the assistant. The interface also treats automation-level changes as interactive transitions rather than silent mode switches: the current level of automation remains continuously visible, transitions are signaled, and the ATCO can intervene at any time. Where explanatory functionality is available, explanations are designed to be lightweight and operationally timed (e.g., concise rationales that answer what changed, why, and what happens next) to support trust calibration without interrupting the primary control activity.

In CODA, explanation is therefore not conceived as a generic explainable AI layer exposing the internal structure of the adaptive algorithm. It is conceived as an event-based operational rationale attached to delegation-relevant changes. The design question is not simply *why did the automation level change?*, but rather *what changed in my work situation, which task or flight is affected, why is CODA proposing or initiating this delegation now, what will happen if I do nothing, and what control action remains available to me?*. These questions were derived from workshop feedback emphasising the need to see what CODA is doing, what each automation state means, and how to retain authority without creating a parallel monitoring task. Accordingly, explanation content is deliberately selective: it exposes the local trigger and consequence of an adaptive action, not a complete technical account of the prediction model (see Table 3).

Another recurrent design issue concerned whether CODA should be integrated into existing controller working-position displays or presented as a dedicated assistant interface. The revised design rationale does not treat this as a binary choice. For an adaptive delegation system, the preferred operational direction is a layered integration strategy. Primary coordination cues (current automation mode, task ownership, affected flight or task, pending delegation, and immediate control actions such as approve, veto, or reclaim) should be displayed as close as possible to the operational object of interest, for example in or near the radar display, electronic flight strips, or the task timeline. This reduces interaction switching costs and helps the ATCO interpret CODA actions within the current traffic context.

However, full integration also has drawbacks. If assistant-generated information is blended too deeply into certified ATC displays, the agency and status of the assistant may become ambiguous: the ATCO may have difficulty distinguishing controller-owned information, conventional system support, and CODA-generated recommendations. Deep integration may also increase display clutter, create unwanted associations with safety-critical functions, and complicate certification and failure-management requirements. Conversely, a fully standalone assistant display makes CODA's agency and scope more explicit, but risks fragmenting attention and creating an additional monitoring task. The CODA prototype therefore used a dedicated

Table 3 Operational questions addressed by CODA explanation cues. Explanations are treated as selective rationale cues attached to delegation events, not as full technical explanations of the adaptive algorithm

ATCO question	Explanation content	Typical trigger	Interaction mode
What changed?	Mode transition, task-ownership change, pending delegation, or change in task/traffic/readiness state	New recommendation, automatic assignment in *Proactive-Strong*, or change in eligibility state	Minimal automatic notification
Which flight or task is affected?	Aircraft/task identifier, current owner, pending owner, and temporal position relative to decision barrier	Flight card moves, task becomes eligible, or delegation is proposed	Displayed on the task card and timeline
Why now?	Local rationale based on the active trigger class: traffic demand, task-load demand, readiness cue if reliable, or task timing	*Proactive-Light* recommendation or *Proactive-Strong* initiation	Short rationale shown in notification; optional drill-down in the explanation panel
What happens next?	Expected next state if no action is taken, e.g., task remains ATCO-owned, moves to CODA after grace period, or requires confirmation	Pending transition or approaching decision barrier	Preview cue and count-down/decision-barrier display
What can I do?	Available control action: approve, reject, veto, reclaim, inspect, or return to Manual	Any delegation proposal or ownership transition	Explicit control affordance on the affected task/flight

interface as a research and implementation compromise within the project technical environment, not as a claim about the preferred final operational architecture.

5.2 HMI concept and prototypes

Design-process chronology. The HMI design followed a three-stage, workshop-driven process. First, the team derived an initial interaction concept from CODA's operational assumptions and from Workshop 1, where participants discussed which routine ATC activities could be considered for support and which boundaries should constrain any AI-based assistance. This led to the first task-board concept organised around task ownership and temporal sequencing. Second, this concept was translated into low- and medium-fidelity mock-ups using Moqups/Figma and reviewed during Workshop 2, which focused on whether the proposed ownership columns, pending-delegation states, automation modes, and transparency cues were understandable and operationally plausible. Third, the design was implemented as a higher-fidelity research prototype and reviewed during Workshop 3, which focused on transition timing, veto/reclaim mechanisms, mental-state display, integration with radar/supervision, and task criticality. Thus, the prototype figures should be read chronologically: early sketches and mock-ups represent the concept-formation stage, implemented screens represent the intermediate prototype-review stage, and the final HMI represents the formative review stage used to consolidate design requirements.

Interfaces development. The interface operationalises CODA's AA by visualizing task delegation across three levels of proactiveness (*Manual*, *Proactive-Light*, *Proactive-Strong*) and by maintaining a persistent representation of task ownership. In the beginning, the HMI was conceptually organized into three columns (*Unassigned*, *ATCO*, *CODA*) that reflect who currently owns each task, complemented by a timeline that provides short-term temporal context (recent past and near-future tasks). This mapping was intentionally simple to support rapid "at-a-glance" comprehension under time pressure.

In *Proactive-Light* mode, CODA proposes task assignments and waits for ATCO confirmation; visually, suggested tasks are placed in a transient pending position between columns, making the request explicit and easy to accept or reject. In *Proactive-Strong* mode, CODA may initiate task-management ownership for eligible non-critical tasks after notifying the ATCO; this transition remains reversible before a defined commit point through veto or reclaim. In *Manual* mode, the ATCO remains the sole allocator of tasks, with interactions designed to be quick and low-friction (i.e., direct assignment to ATCO or CODA).

Design iterations progressed from low-fidelity mock-ups to higher-fidelity interactive prototypes. Initial sketches (Fig. 4, left) were produced to quickly review layout and interaction primitives. Then, subsequent versions (Fig. 4, right) were refined using standard user interface design tools (e.g., Moqups/Figma) and then translated into implementable components (Fig. 5). This pipeline allowed rapid iteration on interaction details (e.g., how to represent pending delegation, how to signal a level change, and how to express ownership clearly) while maintaining consistency across design and implementation. In addition, it includes task grouping interactions (e.g., assigning tasks by flight or by task type) to reduce repetitive actions when multiple items are affected by the same delegation decision.

The final prototype represents the three automation levels through explicit, reversible interaction mechanisms (Fig. 6) and preserves a modified mapping (*Unassigned* column was removed due to feedback from ATCO, see Sect. 6, thus keeping only the *ATCO* and *CODA* columns). Concretely,

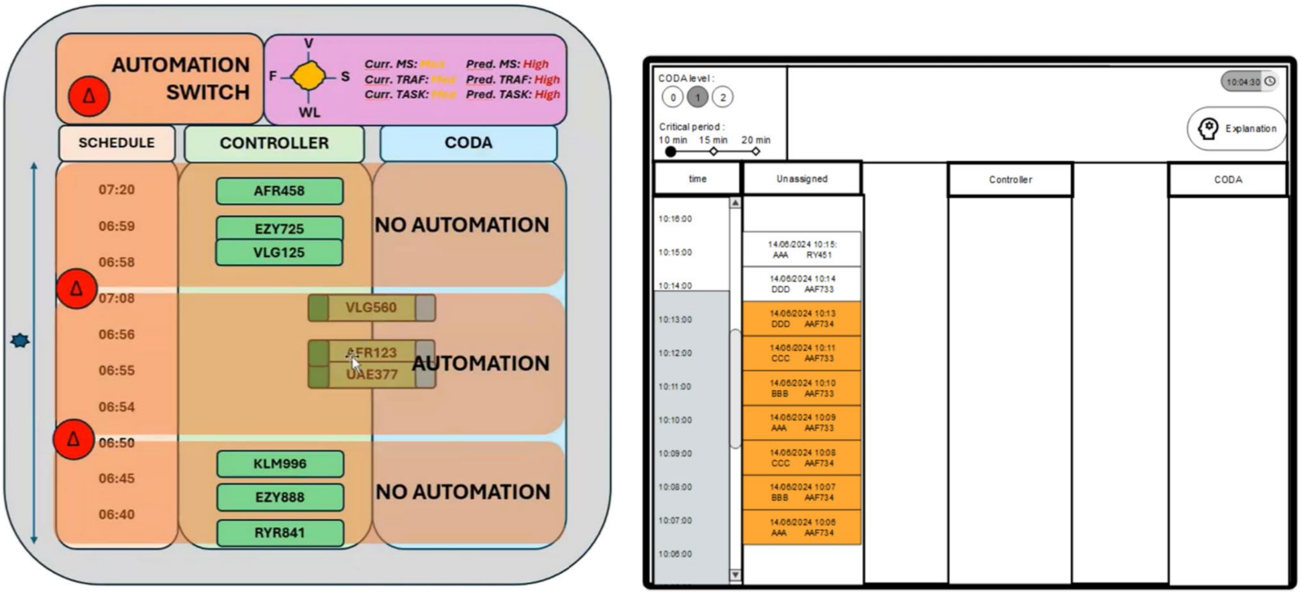

Fig. 4 Left: Early hand-drafted mock-up of the CODA interface illustrating the core task-ownership structure and temporal organization. Right: Examples of higher-fidelity research-prototype screens derived from the mock-ups, used for interactive prototype review and iteration

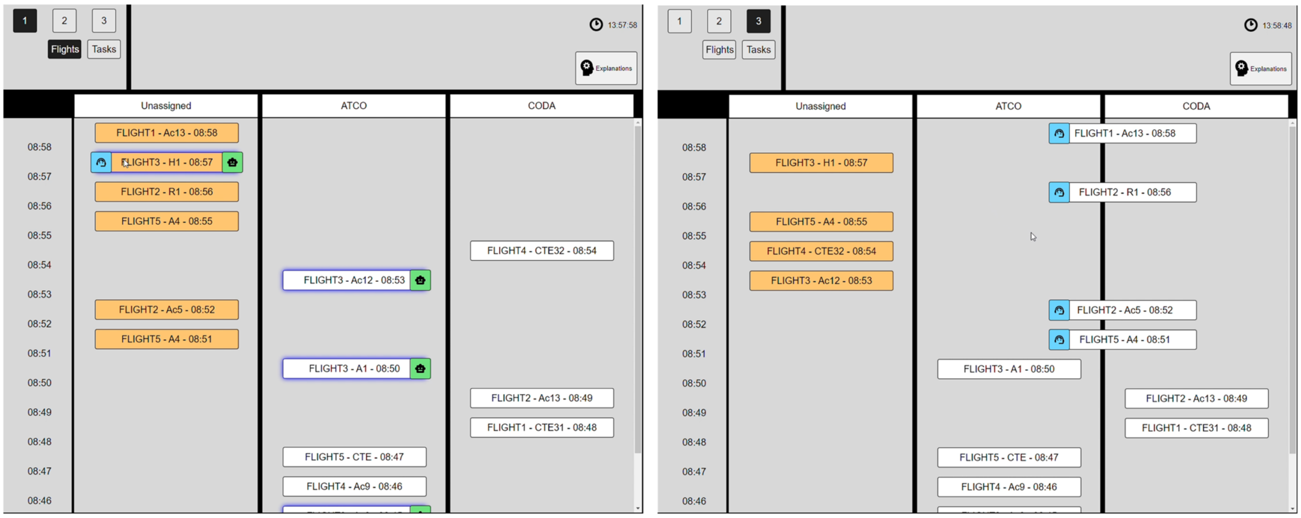

Fig. 5 Examples of higher-fidelity research-prototype screens derived from the mock-ups and used for interactive prototype review and iteration with the expert ATCOs during workshops. Left: In the manual mode, only the ATCO is the allocation agent and the ATCO can assign or leave unassigned the relevant flights and tasks. Right: In the automation mode (Proactive-Strong), the system automatically takes over some flights and tasks and the ATCO can cancel these actions by reassigning the flights and tasks. The blue and green buttons allow the controller to reassign the target flight or task directly to either the controller itself or the CODA system

the interface makes ownership observable as a physical state of each flight card. Aircraft and their associated task items are displayed as movable tokens that transition between columns when responsibility changes, thereby externalizing allocation dynamics and reducing ambiguity during handovers. The HMI also encodes operational salience through a lightweight colour scheme. For example, incoming flights requiring upcoming ATCO management are highlighted (orange), while less time-critical or already-handled items remain in neutral tones, supporting rapid prioritization without adding new symbolic codes that could be confused with those already displayed in the separation-assurance tools.

To support calibrated reliance and maintain supervisory readiness, the final HMI integrates three complementary transparency elements. Once again, the present paper reports these as design requirements and prototype features. It does not claim that the explanation format has yet been shown to improve understanding or trust calibration.

You need to manage flight TAP477Z

OK

Fig. 6 Final CODA HMI for adaptive delegation of non-critical ATC support tasks. Two complementary interfaces are shown: the ATCO view (left), centred on the controller's activity, and the assistant view (right), highlighting CODA-related task management. A vertical timeline on the left supports temporal organisation of upcoming actions, while aircraft labels are distributed across columns reflecting task ownership and automation level, and move between columns as delegation evolves. The horizontal blue lines indicate decision barriers, corresponding to temporal thresholds for task review or execution. The interface was designed to make automation state, task allocation, and transition timing continuously visible as design requirements for transparent and controllable adaptive delegation

First, the prototype includes a two-level explanation mechanism. Minimal cues are displayed automatically when a delegation-relevant event occurs, such as a mode transition, task-ownership change, or *Proactive-Light*/*Proactive-Strong* proposal. More detailed rationale cues are available on demand through the explanation panel. The objective is to support fast verification at the moment of delegation without continuously displaying explanatory material. Its content was specified around operational questions identified during the design process: what changed, why the assistant proposed or initiated a change, which task or flight is affected, and what control options remain available to the ATCO (Fig. 6, Right – green window on the top edge). The active automation mode is made continuously visible through the mode-status area and through the spatial state of task cards. Delegation is primarily item-specific: it applies to an eligible task, flight, or task group displayed in the HMI, not automatically to all aircraft in the sector or to all flights within a time window. Grouping interactions allow the ATCO to apply the same decision to tasks sharing a flight/callsign or task type when appropriate, but this remains an explicit interaction affordance rather than a hidden global mode.

Second, the neuro-adaptive logic is made legible via a compact triangular display that summarizes the three driving components of adaptation (Traffic, Task, Neuro); the current and predicted states are positioned within an "envelope" visualisation (grey and blue colours, respectively), allowing the ATCO to understand at a glance whether adaptation is driven primarily by traffic demand, task pressure, or inferred cognitive readiness (Fig. 6, Left – triangular on the top edge).

Third, automation transitions are never silent: mode changes are explicitly signaled, and salient notifications (e.g., a green banner prompting that a specific aircraft must be managed) provide immediate, actionable cues while preserving the ATCO's authority to accept, override, or veto delegation actions (Fig. 6, Right – green window on the top edge). Together, these design choices instantiate the principle that effective human-automation interaction depends not on maximizing autonomy, but on keeping automation observable, predictable, and recoverable, thereby mitigating out-of-the-loop risks during transitions of control.

Workshops as a design driver and transition to qualitative findings. To make the design trajectory more explicit, Table 4 summarises the chronology of the CODA user-centred design process and relationship between workshops, artefacts, and design decisions. The objective was not to optimise the interface visually at each step, but to progressively align the HMI with three constraints repeatedly identified during the design process: preserving controller authority, making task ownership immediately visible, and limiting transparency to information that supports coordination without increasing workload.

The HMI was not specified as a purely technology-driven artefact; it was iteratively refined through structured user-centred workshops with expert ATCOs. These workshops served to (i) elicit expectations and constraints regarding transparency, authority, and timing of automation transitions, (ii) identify acceptable interaction patterns for delegation and override, and (iii) formatively assess whether CODA's representations were perceived as compatible with ATCO mental models and operational tempo. Section 6 reports the workshop protocol and synthesizes the qualitative themes that directly informed the interface requirements and prototype evolution, ensuring explicit traceability between user input and the resulting HMI design.

6 Workshops: ATCO-driven requirements and design implications

6.1 Qualitative study design, data, and analysis

To align CODA with operational expectations, we conducted an iterative user-centred qualitative study based on three structured workshop sessions with en-route ATCOs at ENAC (Toulouse, France). The main objectives were to (i) elicit operationally relevant, non-critical use cases for AI support, (ii) collect feedback on successive HMI prototypes and interaction logic, and (iii) assess perceived usability, interpretability, and acceptability of CODA's adaptive modes and delegation concept.

Across sessions, the material shown to participants included progressively refined HMI mock-ups and a functional prototype (task timeline + assignment columns), as well as a task taxonomy used to discuss delegation boundaries. Data comprised contemporaneous structured notes and design-iteration artefacts produced during the sessions, including participant questions, critiques, and suggestions captured during demonstrations and guided discussions, and the corresponding decisions taken for subsequent prototype revisions.

We analysed the material using a pragmatic thematic synthesis aligned with the study purpose: we (1) extracted meaning units relevant to delegation boundaries, transparency and mode transitions, mental-state display, integration with existing supervision, training/trust, and task criticality; (2) consolidated recurrent patterns into themes that were stable across sessions; and (3) translated each theme into actionable interface or teaming requirements used to drive subsequent design iterations.

6.2 Participants and sessions

Across the three workshops, 17 unique individuals took part in the sessions, including 13 non-research participants and 4 research-team facilitators. The non-research participants comprised 8 expert en-route ATCO instructors, 3 ATCO students, and 2 engineering/data-visualisation students. All expert ATCOs were ENAC instructors, each with more than 15 years of experience. Across all workshops, the research team acted strictly as facilitators to structure the discussion, probe for clarification, and ensure complete capture of feedback; they did not propose solutions for adoption nor steer participants toward predefined outcomes. To minimize expectation and authority bias, contributions from ATCOs were treated as the primary evidential input, and facilitation was limited to neutral prompts (e.g., requesting examples, rationale, and boundary conditions) and to documenting points of agreement and disagreement. The workshops were structured as follows:

Workshop 1 (Brainstorming - 22 Nov 2023 - 3h). Four expert en-route ATCOs (age: $M = 43.8$, $SD = 1.7$; $n = 4$), three ATCO students (age: $M = 23.3$, $SD = 0.6$; $n = 3$), and the research team (age: $M = 35.7$, $SD = 11.7$; $n = 3$) participated. The session focused on identifying candidate AI-support functions and sketching first interface concepts (Overall - age: $M = 35.2$, $SD = 10.5$, $n = 10$).

Workshop 1 was not framed as an invitation to automate ATC work by default. To avoid an *AI-first* bias, the discussion was structured around current ATCO activities and pain points before any automation function was

Table 4 Chronology of the CODA user-centred design process and relationship between workshops, artefacts, and design decisions

Stage	Material presented	Workshop objective and participant task	Main design question	Resulting design decision
Workshop 1: concept elicitation	CODA project assumptions, en-route task taxonomy, examples of routine/non-critical activities, early paper sketches	Participants were asked to identify routine workflow activities that could be considered for assistance, but also to specify when assistance would be inappropriate. The facilitation explicitly framed AI support as optional and bounded, not as a default objective	Which ATC activities are plausible candidates for support, and which authority/safety boundaries must constrain them?	Focus on non-critical support tasks; initial task-ownership metaphor; first ATCO/CODA task-board concepts; need for visible authority boundaries
Between W1 and W2: mock-up development	Low- and medium-fidelity mock-ups developed in Moqups/Figma	Research-team synthesis of W1 feedback into interaction primitives: ownership columns, pending state, timeline, and three proactivity modes	How can candidate delegation concepts be made visually inspectable and reversible?	Creation of task timeline, ATCO/CODA ownership columns, pending delegation state, approve/reject interaction, and initial mode indicators
Workshop 2: intermediate prototype review	Interactive mock-ups and first implementable screens illustrating task assignment, mode logic, and delegation proposals	Participants reviewed the prototype, commented on usability and operational plausibility, and identified unclear states, risky transitions, or excessive interaction costs	Are the proposed representations of ownership, automation level, and delegation understandable and compatible with ATCO routines?	Confirmation of three-mode logic; refinement of pending delegation representation; emphasis on explicit cancel/override; need to reduce standalone-monitoring burden
Between W2 and W3: high-fidelity implementation	Higher-fidelity research prototype implemented from the design mock-ups within the CODA project technical environment	Research-team integration of W2 feedback into a more stable prototype, including revised ownership display, notification/rationale cues, and automation-state visualisation	Which HMI elements must be stabilised before final formative review?	Removal/refinement of unassigned column; clearer ATCO/CODA ownership columns; explicit transition cues; explanation/notification area; compact adaptation-state display
Workshop 3: final formative review	Final HMI prototype, task-criticality taxonomy, automation-mode examples, and mental-state/adaptation visualisations	Participants reviewed the final concept and were asked to identify acceptance conditions, remaining transparency problems, unsafe delegation cases, and integration requirements	Under which timing, transparency, task-criticality, and integration conditions would CODA-like delegation be acceptable?	Longer/configurable veto window; reclaim until commit point; selective/on-demand explanations; radar/provenance integration requirement; conditional task criticality; training and progressive deployment requirements

considered. Participants were first asked to identify routine activities, coordination burdens, and situations in which additional support could be useful. They were then asked to specify exclusion conditions: tasks that should remain ATCO-owned, contexts in which support would be unsafe or disruptive, and information that would be needed before accepting any delegation. The objective was therefore not to ask ATCOs to design algorithms, but to elicit operational constraints, candidate support opportunities, and boundary conditions that engineers could later translate into AA functions.

Workshop 2 (Intermediate prototype review - 24 Jul 2024 - 2h). Two expert en-route ATCOs (age: $M = 45.5$, $SD = 0.7$; $n = 2$), two engineering students specialised in data visualisation (age: $M = 22.0$, $SD = 0.0$; $n = 2$), and the research team (age: $M = 36.7$, $SD = 11.7$; $n = 3$) reviewed the initial prototypes and discussed delegation logic, transparency, and interaction modes (Overall - age: $M = 35.0$, $SD = 11.8$, $n = 7$).

Workshop 3 (Final formative prototype review - 2 /3 Oct 2024 - 2h) Two expert en-route ATCOs (age: $M = 43.5$, $SD = 0.7$; $n = 2$) and the research team (age: $M = 34.5$, $SD = 10.5$; $n = 4$) reviewed the final HMI and discussed acceptability conditions across seven topics: general HMI design, task assignment representation, adaptive strategy and automation levels, portrayal of mental-state information, integration with radar/supervision, training and trust, and task criticality (Overall - age: $M = 37.5$, $SD = 9.4$, $n = 6$).

6.3 Gathering data

The qualitative material collected during the workshops was analysed through a structured thematic synthesis aimed at translating controller feedback into design requirements (Fig. 7). Given the exploratory and design-oriented nature of the study, the analysis combined deductive and inductive reasoning. A first deductive structure was derived from the main design dimensions of CODA: task delegation, automation-level transitions, transparency and explanation, task ownership, and HMI integration into the controller's workflow. Within this structure, controller comments were reviewed iteratively to identify recurrent concerns, usability issues, operational constraints, and concrete design suggestions. Initial observations were then grouped into higher-level themes and consolidated into design implications for

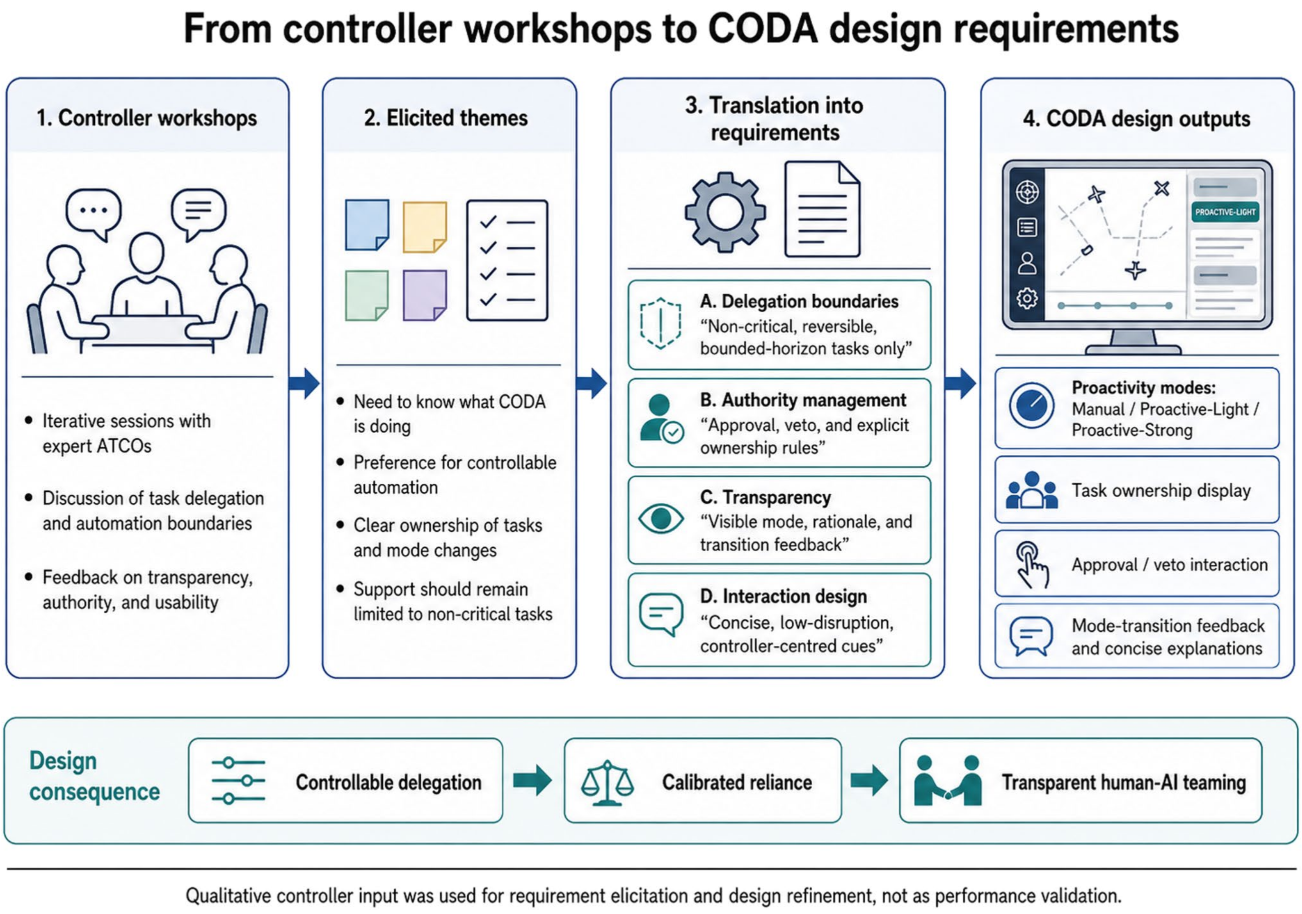
From controller workshops to CODA design requirements

Qualitative controller input was used for requirement elicitation and design refinement, not as performance validation.

Fig. 7 From controller workshops to CODA design requirements. The figure illustrates how qualitative feedback from expert ATCO workshops was translated into operational requirements and HMI design mechanisms for transparent AA. Elicited themes concerning controllability, task ownership, transparency, and the limitation of automation to non-critical tasks were transformed into requirements for delegation boundaries, authority management, transition feedback, and concise interaction cues. These requirements informed the CODA design outputs, including proactivity modes, ownership displays, approval/veto interactions, and mode-transition explanations. The figure frames workshop input as requirement elicitation and design refinement, rather than as performance validation the AA strategy and HMI. The final synthesis was organised around three recurring dimensions: trust and the balance of transparency, preferences for adaptive strategies and automation levels, and HMI mechanisms supporting transparent teaming. These themes were not treated as independent empirical outcomes, but as design-relevant insights used to refine the prototype and clarify the operational boundaries of CODA.

The synthesis was discussed within the research team to ensure that the resulting requirements remained consistent with both controller feedback and the technical constraints of the CODA platform. Iteration decisions were made by combining recurrence, operational criticality, and feasibility. Feedback was prioritised when it was repeated across sessions, when it concerned safety/authority boundaries, when it identified a risk of mode confusion or loss of situation awareness, or when it could be implemented within the prototype constraints. Design changes were therefore not based on frequency alone: a single expert comment could trigger modification when it concerned a safety-critical ambiguity, such as unclear ownership, insufficient veto time, or confusion between CODA support and separation-assurance tools.

6.4 Key themes and findings across workshops

Across the three workshops, ATCOs progressively refined the conditions under which task delegation and adaptive behaviour would be operationally acceptable. Workshop 1 elicited candidate support opportunities and exclusion conditions grounded in current en-route work; Workshop 2 provided an intermediate review of the first task-delegation mock-ups and scrutinised the interaction modes; Workshop 3 provided a final formative review of the HMI concept and articulated concrete acceptance constraints concerning timing, preview horizon, information selectivity, task criticality, and operational integration. These sessions were not treated as performance validation but they were used to transform controller feedback into design requirements.

Theme 1: Delegation is valued for routine tasks, but requires actionable transparency and meaningful labels: Workshop 1 established that delegation is most valuable for routine, workflow-relevant actions that reduce cognitive load without encroaching on conflict-critical decision making. Workshop 2 consolidated these candidate use cases (e.g., straightforward procedural changes and coordination when nominal). Workshop 3 clarified that delegation is only acceptable if CODA's behaviour is inspectable and task codes are understandable: ATCOs noted that "you cannot see what CODA is doing" and that "ATCOs need to know ...what each thing means." This implies (i) drill-down access to task semantics (click/hover on task code), and (ii) short execution-intent descriptions when CODA takes ownership (what will be changed, when, and with which constraints).

Theme 2: Representation of (un)assigned tasks must support prospective control: Workshop 1 explored initial concepts for representing upcoming work and ownership. Workshop 2 supported convergence toward a board-like representation with explicit ownership and intermediate states for pending delegation. Workshop 3 argued that a dedicated "unassigned" column may be cognitively suboptimal, proposing an in-between representation aligned with strip-board practice: "in the middle between the CODA and ATCO columns" and "why are they ...in the separate column?" A design implication is to represent pending tasks as a transient state between columns, and to use visual encoding within the ATCO column to distinguish (i) always-human tasks from (ii) tasks eligible for delegation depending on mode.

Theme 3: Veto/override timing is a workload-sensitive safety mechanism: Workshop 2 highlighted the importance of an explicit cancel/override interaction for proactive delegation (*Proactive-Strong* mode), initially operationalised as a short time window. Workshop 3 strongly challenged the 5 -second veto duration: "5 seconds ...is much too short" and "5 secs is not enough ...10 is better, 15 even better." ATCOs also suggested a state-dependent reclaim possibility: "if CODA hasn't started ...you can still get it back." This supports (i) increasing and/or making configurable the veto window, and (ii) introducing a reclaim rule until a defined commit point (execution start).

Theme 4: Mode changes must be predictable, previewed, and aligned with ATCO planning horizons: Workshop 1 discussed the need for interpretable adaptation rather than opaque automation escalation. Workshop 2 converged on a small number of automation levels (three) as a prerequisite for maintaining a workable mental model. Workshop 3 specified concrete preview and stability requirements: a colour-coded mode representation in the time bar ("colour code ...presented in the time bar") and multi-minute anticipation ("5–6 minutes ...in the past and future"; "impacts ...strategy ...5 minutes ahead"; "three minutes becomes uncomfortable"). The implication is conservative switching (hysteresis/dwell times) plus a multi-minute transition preview embedded in the main temporal representation.

Theme 5: Mental-state information must be conservative and primarily on-demand: Workshop 2 raised a core acceptance question: how the system infers fatigue/workload and how this should (or should not) be surfaced to the ATCO. Workshop 3 refined this into a minimalist principle: ATCOs preferred essential alerts over continuous detail ("only needed 'FATIGUE' ") and warned that too much detail can distract. At the same time, on-demand access was viewed as acceptable (“if I can access it when I want”). This supports (i) threshold-based alerting for safety-relevant conditions, (ii) optional drill-down explanations (measurement provenance and confidence) outside the main task display, and (iii) avoiding high-frequency micro-feedback that could induce maladaptive self-monitoring.

Theme 6: Operational fit requires radar integration, provenance cues, and escalation pathways: Workshop 1 positioned operational integration as a key acceptance determinant (avoid standalone tools that fragment attention). Workshop 2 emphasised that CODA must remain tightly coupled to ATCO routines and should not create additional interaction overhead. Workshop 3 made two requirements explicit: radar coupling (“integration on the radar screen ...helpful”) and provenance marking for CODA actions (“adding a CODA symbol ...next to the parameter ...changed”). ATCOs also proposed escalation to supervision under overload/fatigue conditions (“alert to the supervisor”). Together, these findings motivate clear provenance cues, radar-oriented surfacing of salient events, and a supervisory alert channel for wellbeing/overload triggers (aligned with operational practice).

Theme 7: Trust depends on progressive training, demonstrated benefit under workload, and moderation to protect skills: Workshop 2 framed trust as conditional on efficiency: ATCOs will not adopt a system that requires constant verification. Workshop 3 articulated training and demonstration conditions: “it needs to be trained” and “You don’t want to lose time by checking.” They recommended demonstrating benefit in short, demanding sequences (“5 -10 minutes ...significant traffic”), and warned about long-term deskilling if automation dominates routine work (“moderate use ...for keeping the skills”). This implies staged deployment (low intensity first), evaluation under medium/high workload (non-emergency), and adaptive delegation policies that preserve meaningful engagement to mitigate skill decay.

6.5 Task criticality is context-dependent: delegation rules must combine task type and context

Across workshops, ATCOs repeatedly returned to the principle that CODA should prioritise non-critical support, but that “criticality” is not a fixed property of a task label: it depends on coupling, traffic context, and urgency. Workshop 3 operationalised this by classifying the project task set into critical, not critical, and unsure/conditional categories, with explicit contextual exceptions (e.g., holding tasks depending on the number of aircraft). This supports a conservative delegation policy in which CODA combines task category with contextual risk indicators before proposing or initiating delegation.

Table 5 summarises the cross-workshop evidence and its translation into concrete interface and policy requirements, providing traceability from qualitative insights to design decisions. Figure 8 then clusters these requirements into higher-level recommendations for implementation.

7 Discussion and conclusive remarks

This paper reports the CODA outcomes that specify and operationalise an AA strategy for non-critical ATC support tasks through (i) explicit mode logic and delegation rules, and (ii) an interface-mediated teaming model that keeps ATCO authority observable and enforceable at interaction time. Methodologically, the contribution is design- and requirements-led: we translate AA principles into a concrete delegation policy (three graduated proactive modes) and a set of HMI mechanisms, iteratively refined through structured engagements with expert ATCOs and materialised in successive mock-ups, a high-fidelity prototype, and an implemented interface. The term implemented interface refers here to a research prototype developed within the CODA project environment, not to an operationally integrated or certified ATC system. In this framing, CODA’s primary evidence is not a performance claim, but a defensible specification of “how adaptation is made governable” in a high-tempo supervisory setting, consistent with the view that dynamic automation succeeds only when authority, observability, and transitions are deliberately engineered rather than left implicit (Endsley 2017; Feigh et al. 2012; Kaber and Endsley 2004; Lee and See 2004; Parasuraman et al. 2000; Wickens 1999).

With respect to **RQ1** regarding mode logic and authority contracts, the results indicate that proactive delegation becomes acceptable when it is constrained by graduated authority and reversible commitment. CODA’s *Manual* / *Proactive-Light* / *Proactive-Strong* structure externalises this contract: it differentiates recommendation from initiative, binds each mode to explicit approval/notification conditions, and preserves the ATCO’s primacy through rapid veto/override pathways and persistent cues of who currently owns each delegated task. Across workshops, acceptance repeatedly hinged on the system’s ability to (i) render the next automation action predictable, (ii) keep the scope and consequences of delegation inspectable before commitment, and (iii) support swift reversion when local tactical intent or radio/sector dynamics make delegation undesirable. Delegation may be undesirable even for a normally eligible task when the local operational context makes the task part of a broader tactical plan that is not fully represented

Table 5 Cross-workshop qualitative evidence and resulting CODA requirements (W1: Brainstorming; W2: Intermediate prototype review; W3: Final formative prototype review).

Theme	Evidence across workshops	Design requirement
Delegation boundaries & transparency	Candidate non-critical use cases identified; need to preserve ATCO authority (W1). Review of routine-delegation value; scrutiny of interaction modes (W2). "you cannot see what CODA is doing"; "ATCOs need to know ...what each thing means" (W3).	Task drill-down semantics; execution-intent description; approval/veto/override.
Representation of pending work	Early concepts for task ownership and scheduling (W1). Preference for board-like representation with explicit ownership (W2). "in the middle between the CODA and ATCO columns"; "why ...in the separate column?" (W3).	Pending state between columns; visual encoding for delegable vs. non-delegable tasks.
Veto/override timing	Need for explicit cancel/override for proactive delegation (W2). "5 seconds ...is much too short"; "10 is better, 15 even better"; "if CODA hasn't started ...you can still get it back" (W3).	Longer/configurable veto; reclaim until commit point.
Predictable mode changes	Need for interpretable adaptation (avoid opaque escalation) (W1). Three levels sufficient for workable mental model (W2). "colour code ...presented in the time bar"; "5–6 minutes ...past and future"; "impacts ...strategy ...5 minutes ahead" (W3).	Multi-minute transition preview; hysteresis/dwell times; rate-limited switching.
Mental-state display	Acceptability questions on fatigue/workload inference and display (W2). "only needed 'FATIGUE' " and "if I can access it when I want" (W3).	Threshold-based alerts; on-demand drill-down; separate from primary display.
Operational fit & escalation	Integration and workflow fit emphasised (avoid standalone tools) (W1). Minimise interaction overhead; preserve routines (W2). "integration on the radar screen ...helpful"; "adding a CODA symbol ...next to ...changed"; "alert to the supervisor" (W3).	Radar coupling; provenance cues; supervisor alert channel.
Trust, training & skill retention	Adoption conditional on efficiency (avoid constant checking) (W2). "it needs to be trained"; "You don't want to lose time by checking"; "5-10 minutes ...significant traffic"; "moderate use ...for keeping the skills" (W3).	Progressive deployment; evaluation under workload; preserve engagement to reduce deskilling.

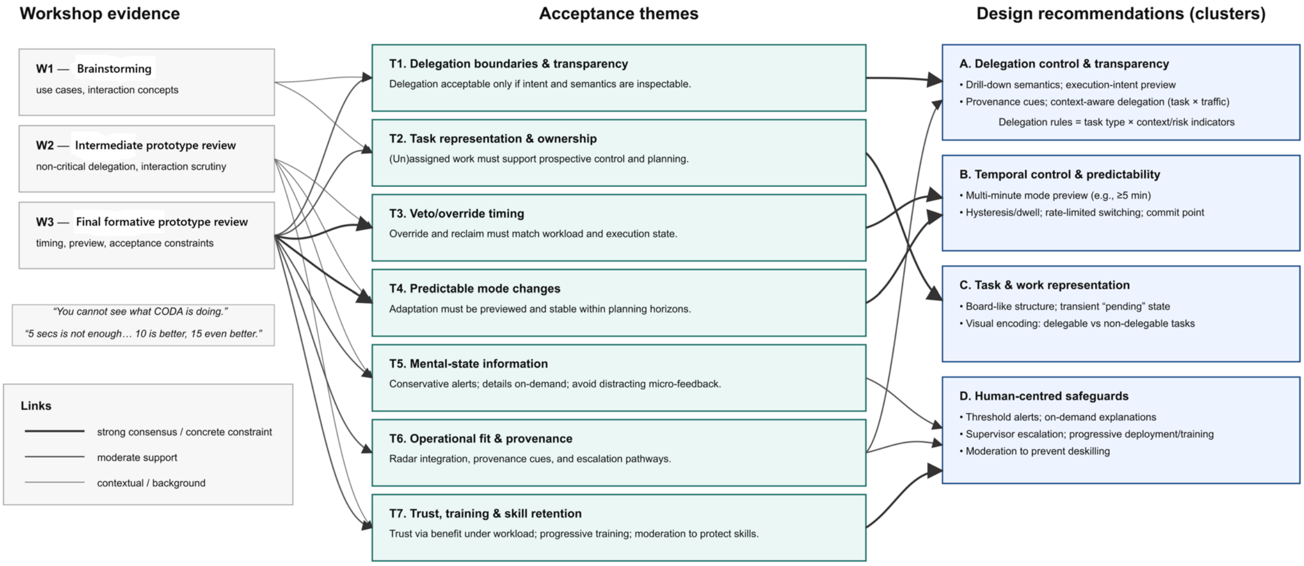

Fig. 8 Synthesis of cross-workshop qualitative findings into actionable CODA design recommendations. Seven acceptance themes were extracted and four design recommendations were identified as clusters

in CODA. This includes situations in which the ATCO is intentionally delaying an instruction to preserve sequencing options, waiting for a pilot readback, coordinating an unresolved transfer condition with an adjacent sector, managing frequency congestion, or anticipating a short-term interaction between aircraft that has not yet become a formal conflict. In such cases, the appropriate ATCO reaction is to reject, defer, veto, or reclaim the task, or to keep the relevant flight in Manual mode. The HMI therefore treats eligibility as conditional: a task may be technically classed as non-critical, but still remain ATCO-owned when tactical intent, communication state, or sector coordination makes delegation unsafe, confusing, or inefficient. These findings align with established principles of appropriate reliance and out-of-the-loop risk mitigation, where stable control transitions and unambiguous authority are central to safe human-automation coordination (Endsley and Kiris 1995; Endsley 2017).

Regarding **RQ2** considering HMI mechanisms to externalise automation state and ownership and support trust calibration with minimal overhead, the workshops converge on a small set of high-value interaction mechanisms: continuous, glanceable mode visibility; explicit task-ownership labelling (e.g., ATCO-owned, pending approval, assistant-owned); previewable delegation proposals with bounded horizons; and concise, actionable rationales anchored in operational triggers (traffic and task-demand indices, and when available human-readiness cues). Critically, participants did not request richer or more frequent explanations; instead, they favoured selective, context-dependent disclosure that supports fast verification at the moment of delegation without creating a parallel monitoring task. This is why explanation in CODA was restricted to delegation-relevant questions rather than designed as a general-purpose model explanation. The HMI must help the controller verify a proposed or ongoing task transfer, not inspect the full internal computation of the adaptive policy. This answers **RQ2** in a principled way: calibrated trust in this setting is better served by timely, low-cost cues that make intent and accountability inspectable at the point of action than by exhaustive transparency that competes with core control activities (Amershi et al. 2019; Cummings 2017; Lee and See 2004; Sheridan and Parasuraman 2005).

Additionally, future CODA-like systems should combine contextual integration for primary delegation cues with peripheral, on-demand access to secondary information. In this sense, the workshops request that explanations should remain *outside the main task display* does not mean that CODA should be disconnected from the ATCO's workflow. Rather, detailed rationales, logs, and diagnostic information should not obscure the primary traffic picture; they should remain available through an adjacent or drill-down layer when the ATCO chooses to inspect them. The expression *outside the main task display* should therefore be understood as a display-layering requirement, not as a requirement for a separate standalone tool. A key risk of full integration is agency dilution: if CODA-generated recommendations are visually indistinguishable from conventional ATC information, the ATCO may misattribute the origin, authority, or reliability of the information. For this reason, integrated CODA cues should include provenance markers, consistent colour/shape coding, and explicit ownership labels distinguishing ATCO-owned, CODA-owned, pending, and unassigned tasks. Such integration would also need to account for the certification burden associated with embedding assistant-generated cues into operational ATC displays.

For **RQ3** about ATCOs constraints on delegation boundaries and transition behaviour, the strongest and most stable outcome is that acceptability is bounded by task criticality and transition discipline. Delegation is valued for routine, workload-relevant support tasks, but becomes unacceptable when boundaries blur with separation assurance responsibilities, when timing conflicts with operational tempo, or when mode changes and handovers are not explicitly previewed and controllable. The thematic synthesis (Fig. 8) therefore supports a general implication for AA in ATC. Proactivity should be treated as a coordination mechanism governed by explicit interaction contracts (preview, approval/veto, ownership, auditability), not as an optimisation objective pursued silently in the background (Langan-Fox et al. 2009; Miller and Parasuraman 2007). In CODA, this implies that the HMI is not a presentational layer but the primary control surface through which adaptive behaviour remains accountable and recoverable under time pressure.

Although CODA is primarily framed as AA, it should not be interpreted as a purely machine-driven allocation system. The adaptive component lies in the computation of mode recommendations or bounded task-management transitions from traffic, task-load, and mental-state/readiness indicators. The adaptable component lies in the ATCO's ability to control these transitions through *Manual* mode selection, approval or rejection of *Proactive-Light* recommendations, veto of *Proactive-Strong* delegation, and reclaim of delegated task-management ownership before the defined commit point. CODA is therefore better described as a hybrid adaptive/adaptable delegation concept (Bernabei and Costantino 2024). The system may initiate support when predefined conditions are met, but human authority remains embedded in the interaction loop through confirmation, cancellation, veto, and reclaim mechanisms. In the prototype, reclaim is represented as a transfer of task-management ownership in the HMI rather than as an operational handover of separation authority. When the ATCO vetoes or reclaims a CODA-managed item, the affected task card moves back to the ATCO-owned state, the pending or active CODA action is cancelled when still before the commit point, and the event is recorded in the notification/task log. If several tasks share the same callsign or task type, grouped reclaim can be represented as an explicit interaction affordance to avoid repeated item-by-item cancellation. After a commit point, reclaim no longer means undoing an action automatically; it means that the ATCO resumes ownership and manages any subsequent correction or follow-up through normal ATC procedures. This hybrid positioning reflects a trade-off. Adaptive triggering can reduce the need for the ATCO to continuously manage support levels during workload peaks and can allow the assistant to anticipate demand changes. However, if adaptation is too autonomous, it may create surprise, mode confusion, or perceived loss of authority. Adaptable control preserves ATCO agency and allows local tactical intent, radio state, and coordination context to override system suggestions, but it can also create additional interaction workload if every support change requires manual management. CODA therefore uses adaptive triggering for bounded proposals or notified transitions, while preserving adaptable control for acceptance, veto, reclaim, and return to *Manual* mode.

Several limitations delimit the claims that can be made from the present evidence and future work follows directly from the research questions and the constraints articulated by ATCOs. First, the main limitation is that the present study does not quantify the effects of the proposed HMI mechanisms on understanding speed, workload, situation awareness, trust calibration, or operational performance. The evidence reported here supports the specification of controller-derived design requirements and prototype features. Demonstrating that these features produce measurable human-performance or operational benefits requires controlled human-in-the-loop validation, which is addressed in companion CODA validation work (Bonelli et al. 2025) and remains a key direction for future research. Second, the workshop evidence is based on a small number of expert ATCOs engaged in iterative design activities. While this is appropriate for requirements elicitation and interaction design refinement, participants may not be representative of the broader European ATCO population, and may exhibit higher-than-average familiarity with, or openness toward, advanced automation concepts. Future work should therefore replicate the elicitation of operational requirements with a larger and more heterogeneous cohort of experts across operational contexts. Third, although task criticality can be context-dependent, CODA is explicitly scoped to non-critical support functions and does not assume separation-critical decision making. Any escalation beyond this boundary is treated as out-of-scope for the present design specification and would require dedicated safety arguments and validation. The proposed logic and interface principles are expected to transfer conceptually, but their direct applicability to conflict management or other safety-critical functions requires dedicated hazard analysis and evidence. The task criticality taxonomy should be refined into operationally testable eligibility criteria for delegation, including explicit boundary conditions and recovery behaviours. Fourth, the adaptive policy accommodates human-state inputs, but the end-to-end sensing pipelines and estimator validation are treated in companion works (Ronca et al. 2024; Giorgi et al. 2025; Dello Iacono et al. 2025; Ricci et al. 2025). Consequently, the present findings speak to how human-state information should be integrated and disclosed (selectively, actionably, and with minimal overhead), rather than to the validity of any specific neurophysiological estimator. Fifth, the adaptive parameters that govern stability (threshold calibration, hysteresis, dwell times, and transition guards) should be specified as auditable configuration artefacts aligned with sector practices and training needs. Sixth, scenario-based evaluations should compare fixed-mode baselines to CODA's interface-mediated adaptation to test the central hypothesis emerging from the workshops: that previewable, vetoable, and ownership-explicit proactivity reduces coordination costs while preserving situation awareness and appropriate reliance. Such studies should quantify not only task outcomes, but also teaming-relevant measures, thereby evaluating AA as a team process rather than a purely algorithmic allocation rule (Endsley 2023; Chen et al. 2025; Lashley et al. 2019; Human-AI Teaming 2022).

While this paper focuses on the adaptive delegation strategies and the HMI, it is highly recommended that future research should also focus on reliably measuring an ATCO's state by means of (neuro)physiological measures such as EEG, ECG, EOG and EDA, which serve as input for AA. While numerous changes in (neuro)physiology have been reported in different mental states, for example increased heart rate or reduced standard deviations of the NN intervals for mental fatigue (Goodman et al. 2025), these parameters can be influenced by numerous other external and internal factors, making it difficult to distinguish the effect of a change in mental state from the effect of other influences. Moreover, the effect of some mental states on (neuro)physiological parameters is sometimes even under discussion, with different studies reporting different effects (Goodman et al. 2025), or some mental states may have a delayed effect on parameters (Giorgi et al. 2023). Hence, more research is needed on ensuring that such parameters can predict a person's state in a reliable and valid manner.

In conclusion, CODA contributes a defensible specification of how AA can be made governable in ATC support settings: a small, workable mode set; explicit authority contracts for delegation; and interface mechanisms that externalise automation state, task ownership, and bounded rationales at the moment of action. In the broader trajectory toward AI-enabled ATM, these design artefacts address a key adoption barrier, trust calibration through inspectable control, by shifting the focus from raw automation capability to interaction-time accountability. This positions CODA as a concrete step toward operationally acceptable HAIT for workload-relevant but non-critical tasks, and a basis for subsequent quantitative validation of performance, safety, and human-state-triggered adaptation in larger-scale evaluations.

Acknowledgements We would like to warmly thank all the expert air traffic controllers, student controllers, and engineers who participated in the various workshops of this project and helped ensure its successful completion.

Author contributions A.V., S.K., G.T., L.M., and C.H. contributed to the conceptualization, methodology, and formal analysis and investigation. A.V. and S.K. wrote the original draft. A.V., S.K., G.T., L.M., and C.H. reviewed and edited the manuscript. C.H. acquired funding, provided resources, and supervised the project. All authors reviewed the manuscript.

Funding Open access funding provided by École Nationale de l'Aviation Civile. This work is part of the CODA Exploratory Research project, funded by the SESAR Joint Undertaking under the Digital European Sky initiative (Grant Agreement No. 101114765).

Data availability No datasets were generated or analysed during the current study.

Declarations

Conflict of interest The authors declare no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

References

Amershi S, Weld D, Vorvoreanu M, Fourney A, Nushi B, Collisson P, Suh J, Iqbal S, Bennett PN, Inkpen K, et al (2019) Guidelines for human-ai interaction. In: Proceedings of the 2019 Chi conference on human factors in computing systems, pp 1–13. https://doi.org/10.1145/3290605.3300233

Andrews RW, Lilly JM, Srivastava D, Feigh KM (2023) The role of shared mental models in human-ai teams: a theoretical review. Theor Issues Ergon Sci 24(2):129–175. https://doi.org/10.1080/1463922X.2022.2061080

Aricò P, Borghini G, Di Flumeri G, Sciaraffa N, Colosimo A, Babiloni F (2017) Passive bci in operational environments: insights, recent advances, and future trends. IEEE Trans Biomed Eng 64(7):1431–1436. https://doi.org/10.1109/TBME.2017.2694856

Aricò P, Sciaraffa N, Babiloni F (2020) Brain-computer interfaces: toward a daily life employment. Brain Sci 10(3):157. https://doi.org/10.3390/brainsci10030157

Aricò P, Borghini G, Di Flumeri G, Colosimo A, Bonelli S, Golfetti A, Pozzi S, Imbert J-P, Granger G, Benhacene R, Babiloni F (2016) Adaptive automation triggered by eeg-based mental workload index: a passive brain-computer interface application in realistic air traffic control environment. Frontiers in Human Neuroscience 10:539 https://doi.org/10.3389/fnhum.2016.00539

Balfe N, Sharples S, Wilson JR (2015) Impact of automation: measurement of performance, workload and behaviour in a complex control environment. Appl Ergon 47:52–64. https://doi.org/10.1016/j.apergo.2014.08.002

Berberian B, Somon B, Sahaï A, Gouraud J (2017) The out-of-the-loop brain: a neuroergonomic approach of the human automation interaction. Annu Rev Control 44:303–315. https://doi.org/10.1016/j.arcontrol.2017.09.010

Bernabei M, Costantino F (2024) Adaptive automation: status of research and future challenges. Robot Comput-Integr Manuf 88:102724. https://doi.org/10.1016/j.rcim.2024.102724

Bolić T, Ravenhill P (2021) Sesar: the past, present, and future of European air traffic management research. Engineering 7(4):448–451. https://doi.org/10.1016/j.eng.2020.08.023

Bonelli S, Vicario A, Levantesi A, Torsi S, Iglesias B, Besada J, Hurter C, Veyrié A, Lasheras RG, Frutos PLD, Borghini G, Cañas J (2025) Human-machine performance envelope: controller adaptive digital assistant evaluation. In: SESAR Innovation Days 2025, SESAR Joint Undertaking, pp 1–9. https://doi.org/10.61009/SID.2025.1.48

Cartocci G, Veyrié A, Cavagnetto N, Hurter C, Degas A, Ferreira A, Ahmed MU, Begum S, Barua S, Inguscio BMS et al (2026) Explainable artificial intelligence in air traffic control: effects of expertise on workload, acceptance, and usage intentions. Brain Informatics. https://doi.org/10.1186/s40708-025-00287-6

Causse M, Imbert J, Duchevet A, Veyrié A, Hurter C (2025) Digital assistant in aviation: monitoring, understanding, and supporting operators. Human Factors in Design Engineering and Computing AHFE Open Access 199: AHFE International https://doi.org/10.54941/ahfe1007030

Chancey ET, Politowicz MS, Ballard KM, Unverricht J, Buck BK, Geuther S (2025) Human-automation trust development as a function of automation exposure, familiarity, and perceived risk: a high-fidelity remotely operated aircraft simulation. J Cognitive Eng Decis Making 19(2):200–222. https://doi.org/10.1177/15553434241296573

Chen C, Liao M, Sundar SS (2024) When to explain? Exploring the effects of explanation timing on user perceptions and trust in AI systems. In: Proceedings of the second international symposium on trustworthy autonomous systems, pp 1–17. https://doi.org/10.1145/3686038.3686066

Chen Z, Luo Y, Sra M (2025) Engaging with ai: How interface design shapes human-AI collaboration in high-stakes decision-making. https://doi.org/10.48550/arXiv.2501.16627

Chen D, Yoon HJ, Wan Z, Alluru N, Lee SW, He R, Moore TJ, Nelson FF, Yoon S, Lim H et al (2025) Advancing human-machine teaming: concepts, challenges, and applications. https://doi.org/10.48550/arXiv.2503.16518

Cocchioni M, Bonelli S, Westin C, Ferreira A, Cavagnetto N (2023) Guidelines for artificial intelligence in air traffic management: a contribution to easa strategy. In: AHFE (2023) international conference. https://doi.org/10.54941/ahfe1003008

CODA Consortium: D5.1 – adaptation and human ai interaction strategy and teaming playbook. Technical Report Deliverable D5.1, Grant Agreement 101114765, SESAR 3 Joint Undertaking (Horizon Europe) (2025). Edition 00.02.00, 04 February 2025, Public (PU)

Cummings ML (2017) Automation bias in intelligent time critical decision support systems. In: Decision Making in Aviation, 289–294. Routledge. https://doi.org/10.2514/6.2004-6313

Damacharla P, Javaid AY, Gallimore JJ, Devabhaktuni VK (2018) Common metrics to benchmark human-machine teams (HMT):

a review. IEEE Access 6:38637–38655. https://doi.org/10.1109/ACCESS.2018.2853560

Degas A, Islam MR, Hurter C, Barua S, Rahman H, Poudel M, Ruscio D, Ahmed MU, Begum S, Rahman MA, Bonelli S, Cartocci G, Di Flumeri G, Borghini G, Babiloni F, Aricò P (2022) A survey on artificial intelligence (ai) and explainable ai in air traffic management: current trends and development with future research trajectory. Appl Sci 12(3):1295. https://doi.org/10.3390/app12031295

Dehais F, Lafont A, Roy R, Fairclough S (2020) A neuroergonomics approach to mental workload, engagement and human performance. Front Neurosci 14:268. https://doi.org/10.3389/fnins.2020.00268

Dello Iacono F, Guinti L, Cecchetti M, Giorgi A, Rossi D, Ronca V, Vozzi A, Capotorto R, Babiloni F, Aricò P et al (2025) Analysis of neurophysiological correlates of mental fatigue in both monotonous and demanding driving conditions. Brain Sci 15(9):1001. https://doi.org/10.3390/brainsci15091001

Di Flumeri G, De Crescenzio F, Berberian B, Ohneiser O, Kramer J, Aricò P, Borghini G, Babiloni F, Bagassi S, Piastra S (2019) Brain–computer interface-based adaptive automation to prevent out-of-the-loop phenomenon in air traffic controllers dealing with highly automated systems. Front Hum Neurosci 13:296 https://doi.org/10.3389/fnhum.2019.00296

Endsley MR (2017) From here to autonomy: lessons learned from human-automation research. Hum Factors 59(1):5–27. https://doi.org/10.1177/0018720816681350

Endsley MR (2023) Supporting human-ai teams: transparency, explainability, and situation awareness. Comput Hum Behav 140:107574. https://doi.org/10.1016/j.chb.2022.107574

Endsley MR, Kiris EO (1995) The out-of-the-loop performance problem and level of control in automation. Hum Factors 37(2):381–394. https://doi.org/10.1518/001872095779064555

EUROCONTROL (2022) Aviation Outlook 2050: Air Traffic Forecast Shows Aviation Pathway to Net Zero CO^2 Emissions. Published: 2022-06-17. Accessed: 20 Jan 2026. https://www.eurocontrol.int/article/aviation-outlook-2050-air-traffic-forecast-shows-aviation-pathway-net-zero-co2-emissions/

EUROCONTROL (2024) Forecast Update 2024–2030: Autumn 2024 Update. https://www.eurocontrol.int/publication/eurocontrol-forecast-2024-2030-autumn-update/. Published: 2024-10-15. Accessed: 20 Jan 2026

European Union Aviation Safety Agency (EASA) (2024) EASA Artificial Intelligence (AI) Concept Paper Issue 2: Guidance for Level 1 & 2 machine learning applications. General publication. Accessed: 19 Jul 2024. https://www.easa.europa.eu/en/document-library/general-publications/easa-artificial-intelligence-concept-paper-issue-2

Feigh KM, Dorneich MC, Hayes CC (2012) Toward a characterization of adaptive systems: a framework for researchers and system designers. Hum Factors 54(6):1008–1024. https://doi.org/10.1177/0018720812443983

Gerdes I, Jameel M, Materne LJ, Bruder C (2025) Synergies in the skies: situation awareness and shared mental model in digital-human air traffic control teams. Aerospace. https://doi.org/10.3390/aerospace12060472

Gerdes I, Jameel M, Hunger R, Christoffels L, Gürlük H (2022) The automation evolves: concept for a highly automated controller working position. In: 33rd congress of the international council of the aeronautical sciences, ICAS 2022, ICAS

Giorgi A, Ronca V, Capotorto R, Vozzi A, Rossi D, Aricò P, Borghini G, Van Gasteren M, Melus J, Petrelli M et al (2025) Beyond the time-on-task: an eeg-driven approach for effective physiological assessment of mental fatigue in simulated and real driving. Front Bioeng Biotechnol 13:1682103. https://doi.org/10.3389/fbioe.2025.1682103

Giorgi A, Ronca V, Vozzi A, Aricò P, Borghini G, Capotorto R, Tamborra L, Simonetti I, Sportiello S, Petrelli M et al (2023) Neurophysiological mental fatigue assessment for developing user-centered artificial intelligence as a solution for autonomous driving. Front Neurorobot 17:1240933 https://doi.org/10.3389/fnbot.2023.1240933

Goodman SP, Collins B, Shorter K, Moreland AT, Papic C, Hamlin AS, Kassman B, Marino FE (2025) Approaches to inducing mental fatigue: a systematic review and meta-analysis of (neuro) physiologic indices. Behav Res Methods 57(4):102. https://doi.org/10.3758/s13428-025-02620-7

Gutiérrez Teuler G, Arnaldo Valdés RM, Gómez Comendador VF, Frutos PM, Rodríguez Rodríguez R (2022) Study of the impact of traffic flows on the atc actions. Aerospace 9(8):467. https://doi.org/10.3390/aerospace9080467

Hancock PA (2022) Avoiding adverse autonomous agent actions. Human-Computer Interaction 37(3):211–236. https://doi.org/10.1080/07370024.2021.1970556

Hoffman RR, Mueller ST, Klein G, Litman J (2018) Metrics for explainable AI: challenges and prospects. arXiv:1812.04608

Human-AI Teaming (2022) State-of-the-art and research needs. The National Academies Press. National Academies of Sciences, Engineering, and Medicine, Washington, DC, USA https://doi.org/10.17226/26355

Hunger R, Christoffels L, Friedrich M, Jameel M, Pick A, Gerdes I, Nahmer PM, Sobotzki F (2024) Lesson learned: design and perception of single controller operations support tools. In: International conference on human-computer interaction, Springer, pp 15–33. https://doi.org/10.1007/978-3-031-60731-8_2

Hurter C, Veyrié A, Kebir S, Truong G, Bonelli S, Borghini G, Aricò P, Babiloni F, Baumgartner M, Besada JA et al (2026) Toward a neurophysiological approach to assess optimal human-machine teaming in the critical environment of air traffic control. In: International conference on human-computer Interaction, Springer, pp 181–199. https://doi.org/10.1007/978-3-032-12392-3_12

Inagaki T (2003) Adaptive automation: design of authority for system safety. IFAC Proc Vol 36(14):13–22. https://doi.org/10.1016/S1474-6670(17)32389-3

Islam MR, Ahmed MU, Barua S, Begum S (2022) A systematic review of explainable artificial intelligence in terms of different application domains and tasks. Appl Sci 12(3):1353. https://doi.org/10.3390/app12031353

Jameel M, Tyburzy L, Gerdes I, Pick A, Hunger R, Christoffels L (2023) Enabling digital air traffic controller assistant through human-autonomy teaming design. In: 2023 IEEE/AIAA 42nd digital avionics systems conference (DASC), IEEE, pp 1–9. https://doi.org/10.1109/DASC58513.2023.10311220

Kaber D, Endsley M (2004) The effects of level of automation and adaptive automation on human performance, situation awareness and workload in a dynamic control task. Theor Issues Ergon Sci 5:113–153. https://doi.org/10.1080/1463922021000054335

Kirwan B (2024) The impact of artificial intelligence on future aviation safety culture. Future Transp 4(2):349–379. https://doi.org/10.3390/futuretransp4020018

Kirwan B (2025) Human factors requirements for human-ai teaming in aviation. Future Transp. https://doi.org/10.3390/futuretransp5020042

Langan-Fox J, Canty JM, Sankey MJ (2009) Human-automation teams and adaptable control for future air traffic management. Int J Ind Ergon 39(5):894–903. https://doi.org/10.1016/j.ergon.2009.04.002

Lashley H, Thorpe A, Tylor R, Grabham A (2019) Measuring effectiveness of human autonomy teaming. (NATO STO HFM-300)

Lee JD, See KA (2004) Trust in automation: designing for appropriate reliance. Hum Factors 46(1):50–80. https://doi.org/10.1518/hfes.46.1.50_30392

Li Q, Ng K, Fan Z, Yuan X, Liu H, Bu L (2021) A human-centred approach based on functional near-infrared spectroscopy for adaptive decision-making in the air traffic control environment: a case study. Adv Eng Informatics 49:101325. https://doi.org/10.1016/j.aei.2021.101325

Luciani DT, Löwgren J, Lundberg J (2020) Designing fine-grained interactions for automation in air traffic control. Cognit Technol Work 22:685–701. https://doi.org/10.1007/s10111-019-00598-9

Malakis S, Baumgartner M, Berzina N, Laursen T, Smoker A, Poti A, Fabris G, Velotto S, Scala M, Kontogiannis T (2023) A framework for supporting adaptive human-ai teaming in air traffic control. In: International conference on human-computer interaction, Springer, pp 320–330. https://doi.org/10.1007/978-3-031-35389-5_22

Marsman L, Lingam SN, Klein Obbink B, Vlasblom J, Zon R, Miltenburg M (2024) Controller adaptive digital assistant (CODA): interim conclusions on the development of a new atc system. In: International conference on intelligent human computer interaction, Springer, pp 254–260. https://doi.org/10.1007/978-3-031-88705-5_21

Metzger U, Parasuraman R (2005) Automation in future air traffic management: effects of decision aid reliability on controller performance and mental workload. Hum Factors 47(1):35–49. https://doi.org/10.1518/0018720053653802

Miller CA, Parasuraman R (2007) Designing for flexible interaction between humans and automation: delegation interfaces for supervisory control. Hum Factors 49(1):57–75. https://doi.org/10.1518/001872007779598037

Munoz-de-Escalona E, Leva MC, Gianini G, Frutos PL, Jadronova M (2025) Air traffic controllers communication analysis as a proxy of task demand and mental workload: using voice recording markers for safety critical task. Saf Sci 191:106928. https://doi.org/10.1016/j.cognition.2007.09.003

Nylin M (2023) Flexible automation in air traffic control through adaptation of human-automation collaboration. Linköping Studies in Science and Technology Doctoral Dissertations https://doi.org/10.3384/9789180753371

Ortner P, Steinhöfler R, Leitgeb E, Flühr H (2022) Augmented air traffic control system–artificial intelligence as digital assistance system to predict air traffic conflicts. AI 3(3):623–644. https://doi.org/10.3390/ai3030036

Pacherie E (2008) The phenomenology of action: a conceptual framework. Cognition 107(1):179–217. https://doi.org/10.1016/j.cognition.2007.09.003

Parasuraman R (2000) Designing automation for human use: empirical studies and quantitative models. Ergonomics 43(7):931–951. https://doi.org/10.1080/001401300409125

Parasuraman R, Sheridan TB, Wickens CD (2000) A model for types and levels of human interaction with automation. IEEE Tran Syst Man Cybern- Part A Syst Humans 30(3):286–297. https://doi.org/10.1109/3468.844354

Pham DT, Ali H, Fennedy K, Hsieh MH, Alam S, Duong V (2024) Human-ai hybrid paradigm for collaborative air traffic management systems. In: Proceedings of the SESAR innovation days 2024. Rome, Italy (Paper 087)

Ricci A, Ronca V, Capotorto R, Giorgi A, Vozzi A, Germano D, Borghini G, Di Flumeri G, Babiloni F, Aricò P (2025) Understanding the unexplored: a review on the gap in human factors characterization for industry 5.0. Appl Sci 15(4):1822. https://doi.org/10.3390/app15041822

Ronca V, Capotorto R, Di Flumeri G, Giorgi A, Vozzi A, Germano D, Virgilio VD, Borghini G, Cartocci G, Rossi D et al (2024) Optimizing eeg signal integrity: a comprehensive guide to ocular artifact correction. Bioengineering 11(10):1018. https://doi.org/10.3390/bioengineering11101018

Rooij G, Tisza A, Borst C (2024) Flight-based control allocation: towards human–autonomy teaming in air traffic control. Aerospace. https://doi.org/10.3390/aerospace11110919

Sahaï A, Desantis A, Grynszpan O, Pacherie E, Berberian B (2019) Action co-representation and the sense of agency during a joint simon task: comparing human and machine co-agents. Conscious Cogn 67:44–55. https://doi.org/10.1016/j.concog.2018.11.008

Sahaï A, Caspar E, De Beir A, Grynszpan O, Pacherie E, Berberian B (2023) Modulations of one's sense of agency during human-machine interactions: a behavioural study using a full humanoid robot. Q J Exp Psychol 76(3):606–620. https://doi.org/10.1177/17470218221095841

Scerbo MW (2018) Theoretical perspectives on adaptive automation. In: Automation and human performance, CRC Press, pp 37–63

Sebastiani M, Di Flumeri G, Aricò P, Sciaraffa N, Babiloni F, Borghini G (2020) Neurophysiological vigilance characterisation and assessment: laboratory and realistic validations involving professional air traffic controllers. Brain Sci 10(1):48. https://doi.org/10.3390/brainsci10010048

SESAR Joint undertaking: European ATM Master Plan. https://www.sesarju.eu/masterplan/, Brussels, Belgium (2025)

Sheridan TB, Parasuraman R (2005) Human-automation interaction. Rev Hum Factors Ergon 1(1):89–129. https://doi.org/10.1518/155723405783703082

Svensson, Lundberg J, Forsell C, Rönnberg N (2021) Automation, teamwork, and the feared loss of safety: air traffic controllers' experiences and expectations on current and future atm systems. In: Proceedings of the 32nd European conference on cognitive ergonomics. https://doi.org/10.1145/3452853.3452855

Van Droogenbroeck C, Rankova E, Papenfuss A, Bos T, Zon R (2025) Human-ai teaming–challenges from a practitioner's perspective. In: International conference on human-computer interaction, Springer, pp 337–354. https://doi.org/10.1007/978-3-031-93718-7_22

Wickens CD (1999) Automation in air traffic control: the human performance issues. Automation technology and human performance: current research and trends. pp 2–10

Xie Y, Pongsakornsathien N, Gardi A, Sabatini R (2021) Explanation of machine-learning solutions in air-traffic management. Aerospace 8(8):224. https://doi.org/10.3390/aerospace8080224

Zamarreño Suárez M, Arnaldo Valdés RM, Pérez Moreno F, Delgado-Aguilera Jurado R, Frutos PM, Gómez Comendador VF (2022) How much workload is workload? A human neurophysiological and affective-cognitive performance measurement methodology for atcos. Aircr Eng Aerosp Technol 94(9):1525–1536. https://doi.org/10.1108/AEAT-11-2021-0328

Zamarreño Suárez M, Arnaldo Valdés RM, Pérez Moreno F, Delgado-Aguilera Jurado R, Frutos PM, Gómez Comendador VF (2023) Methodology for determining the event-based taskload of an air traffic controller using real-time simulations. Aerospace 10(2):97. https://doi.org/10.3390/aerospace10020097

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.