Final Deliverable + HC Explanations + Process Description

Minerva University

FE51: Cornerstone Civic Project

Prof. Tambasco

April 25, 2026

Final Deliverable:

STRATEGY DOCUMENT

Adobe Life

Workplace AI Strategy Handbook

REFINED CHALLENGE QUESTION

How might we redesign adaptive workplace AI within the Adobe Life App to reduce ***cognitive overload?***

Prepared for	Prepared by
Julia Ip (Product Manager, Adobe Life)	Team Fivefold (Sky Lee, Sofiya Fishman, Sua Park, Lynn Kim, Shehryar Hassan)

This handbook presents three evidence-based strategies for transitioning Adobe Life from a broadcast information platform to an intelligent, adaptive workplace AI system. Each strategy is analyzed across problem definition, mechanism design, evidence base, and expected outcomes.

Visual Demo: https://9qjl8h.csb.app/

EXECUTIVE SUMMARY

Adobe Life currently operates as a broadcast model feed: information is pushed uniformly to all employees regardless of role or location, in other words, untailored. This architecture is no longer tenable. The hybrid and remote workplace has fragmented employee attention across multiple asynchronous channels (e.g., Slack, Outlook, Teams, calendars), generating a compounding cognitive burden that empirical research links directly to productivity loss.

Empirical evidence is clear: digital communication overload correlates with employee burnout at r = 0.96 (Iqbal Soomro et al., 2025). Employees are interrupted every fifteen minutes across digital channels (Unily, 2025), and each interruption requires an average of twenty-three minutes of recovery time before focused work can resume (SAP Concur, 2026). The cumulative effect is a structurally degraded working environment that undermines Adobe's investment in human capital.

This handbook presents **three** strategies for redesigning Adobe Life around adaptive workplace artificial intelligence. Each strategy targets a specific user segment and interaction context. Together, they aim to increase workplace productivity, grounded in *Cognitive Load Theory* (Sweller, 1988). The *Computers Are Social Actors* paradigm (Nass & Moon, 2000), and all three operate within the compliance boundaries of the *EU AI Act Article 5.*

KEY EVIDENCE BASE

Empirical Finding	Source & Implication
r = 0.96 correlation: digital communication overload → employee burnout	Iqbal Soomro et al. (2025). Magnitude warrants intervention; the unfiltered broadcast model directly endangers staff wellbeing.

Employees are interrupted every 15 minutes by digital notifications	Unily (2025). The platform must reduce notification volume, not increase it.
23 minutes average recovery time per interruption	SAP Concur (2026). Cost per interruption is severe; total daily productivity loss is substantial.
100+ emails and 150+ messages received daily per employee	Microsoft Work Trend Index (2025). Manual triage is cognitively unsustainable at scale.
Three simultaneous demands at office re-entry create peak cognitive load	Sweller, Ayres & Kalyuga (2025). Re-entry design must collapse parallel tasks into a single intent.
Social responsiveness in AI increases trust and adoption	Nass & Moon (2000). Design implications: systems must exhibit contextual, human-like interaction patterns.

STRATEGY 01: SMART RE-ENTRY PORTAL

***Target User**: Discontinuous Employees* | ***Context**: Office Re-entry After ≥ 10 Day Absence*

1.1 Problem Definition and User Segment

Discontinuous users are hybrid or remote employees who attend the physical office only a few times per month. Each return visit triggers a specific cognitive crisis: the simultaneous need to process accumulated updates or updates at Adobe.

Cognitive Load Theory distinguishes between *intrinsic load* (the inherent complexity of a task), *extraneous load* (the cognitive cost of poor design), and *germane load* (the cognitive effort that builds competence). The re-entry problem is an *extraneous load* resulting from cognitive cost due to the absence of a consolidated interface (Sweller, Ayres & Kalyuga, 2025). The three demands of re-entry are not individually complex; they are cognitively costly precisely because they must be resolved simultaneously, in an unfamiliar context, with no decision support.

Research on assisted task completion confirms that *reducing* the number of decision points during high cognitive demand significantly improves user performance and confidence (Google DeepMind, 2024). The intervention is therefore to collapse their sequential resolution into a single declared action.

1.2 Core Mechanism: Geofence-Triggered Single Intent

The mechanism is Geofence-Triggered Activation: when the app detects that the user has entered the office perimeter via dual-signal confirmation (GPS and Wi-Fi), the interface transforms into a single prominent action card with a single choice: "Start My Office Day." This design draws directly on Nass and Moon's (2000) CASA paruserslíminating the need for the user to navigate to multiple features under cognitive strain.

Dual-signal confirmation (GPS + Wi-Fi) is a critical design to eliminate mistakes. Choosing a single-signal confirmation activation is susceptible to false positives, such as a user living near the office or passing by the office, triggering the portal unintentionally. Requiring both signals in confirmation before activating the re-entry mode eliminates this risk.

DESIGN PRINCIPLE: THE SINGLE INTENT

Assisted task-completion research demonstrates that the optimal intervention during high-cognitive-load moments is to provide fewer options. The 'Start My Office Day' card presents exactly **one** decision point.

1.2.1 THREE PARALLEL BACKEND PROCESSES

Upon a single click, the system simultaneously executes three backend processes. The parallelism is architecturally significant: re-entry friction is not caused by any single gap, but by the simultaneous demands of catching up, finding a space, and navigating the environment. Executing them in parallel respects that all three are equally time-critical at the moment of entry.

Process	Mechanism	Specificity
Catch-Up Summary	Scans Slack, Teams, and Outlook for all significant updates since last visit	Surfaces exactly three critical insights, e.g., 'Your team is in Location B today.' Volume is capped at three to prevent overloading.
One-Tap Workplace suggestion	Pattern recognition on historical opt-in data only	**Suggests** a workspace near the user's primary collaborators. All data anonymized. Strictly opt-in. No inference from biometric or location

		history without explicit consent (Compliant with EU AI Act Article 5).
Spatial Wayfinding	AR-lite indoor map via Indoor Positioning Systems (IPS)	Activates **immediately** upon booking confirmation. Guides the user from the lobby to the assigned desk and upcoming meeting rooms. Removes navigational uncertainty from the first minutes of the day. It can be turned off in settings, which is most helpful for newcomers or users who are highly discontinuous.

1.3 EU AI Act Compliance Architecture

Article 5 of the EU AI Act prohibits AI systems from inferring emotional states from biometric data in workplace environments (European Commission, 2026). The Smart Re-entry Portal is designed to operate entirely within this boundary. The geofence trigger relies on location signals (GPS, Wi-Fi) that the user has explicitly granted the app permission to access. The workplace booking suggestion is derived from historical booking patterns and calendar data (both declared behavioral records) and is activated only when the user has opted in to pattern-based suggestions.

Critically, the system does not infer the user's internal state at any point. It does not ask whether the user is stressed, fatigued, or disoriented. It detects the objective condition of office re-entry after a defined absence threshold (ten or more days) and responds with a pre-configured set of support actions.

1.4 Expected Outcomes and Measurement

Outcome	Measurement Method	Baseline Required
Reduction in time to first task completion on re-entry days	timestamp from app open to first task action	Establish a pre-intervention baseline over 60 days
Improvement in belonging scores (remote/hybrid staff)	Quarterly employee engagement survey, segment by user type	Existing survey data; disaggregate discontinuous user cohort
Increased free workspace utilization accuracy	Delta between AI-predicted occupancy and actual data	Current manual booking accuracy rate
Reduction in re-entry support requests	Helpdesk use, questionnaire emails in re-entry context	60-day pre-intervention analysis

VISUAL SIMULATION

STRATEGY 02: CONTEXT FEED

Target User: *Continuous Employees* | ***Context:*** *Daily Working Session, Morning Through End of Day*

2.1 Problem Definition and User Segment

Continuous users engage with Adobe Life every working day. The problem they face is qualitatively different from that of discontinuous users: attrition due to exposure to an unfiltered flood of information. Emails, Teams messages, calendar notifications, and internal updates arrive continuously throughout the day, requiring employees to manually determine which items are urgent, which are important, and which can be safely deferred – a triage task that repeats hundreds of times a day.

The cognitive cost is compounding over time. Mark, Gudith, and Klocke (2008) demonstrate that frequent task-switching and interruptions measurably reduce productivity and increase perceived stress. Microsoft's 2025 Work Trend Index reports that employees receive an average of more than 100 emails and 150+ messages daily. The challenge is not a lack of access to information; users are drowning in it. The challenge is the absence of intelligent pre-processing that converts raw information into structured, actionable guidance before it reaches the employee.

The intervention is not a simple notification filter. It is a system that performs the triage, prioritization, and translation tasks that employees currently perform manually, and surfaces only the output (a structured, actionable briefing) rather than dispersed information.

2.2 Core AI Solution: Three-Layer Processing Architecture

The Context Feed aggregates data from all active channels (Outlook, Microsoft Teams, Calendar, workplace systems, and office presence data) and applies three sequential processing layers before presenting any content to the user. Each layer reduces cognitive demand at a different level:

Processing Layer	Function	Cognitive Load Rationale (Sweller, 1988)
Layer 1: Filter	Removes irrelevant inputs (general newsletters, automated system notifications, low-signal updates) before they enter the user's awareness	Eliminates extraneous cognitive load at source
Layer 2: Prioritize	Ranks remaining content by urgency and importance using task metadata, deadlines, sender signals, and recency	Reduces the decision cost of determining what to address first; priority is pre-determined
Layer 3: Translate	Converts prioritized inputs into clear, one-line next-action statements ['Respond to Alex re: client proposal (Email, Apr 22)'] rather than raw messages	Reduces interpretive cognitive work; the user receives a directive

2.3 Four Intelligence Features

2.3.1 ACTION INTELLIGENCE (TASK DETECTION)

Rather than displaying full messages, the system surfaces only the actionable tasks embedded in each communication. 'Respond to Alex regarding new client proposal (Email, April 22)' replaces a full email thread. 'Confirm meeting with Marina (Teams message)' replaces a message thread. This design prevents the most common productivity failure mode: *emails that were read, noted, and then* ***forgotten*** – leaving the user with the discomfort of unresolved tasks without the reminder to act (Mark, Gudith & Klocke, 2008).

2.3.2 PRIORITY STRUCTURING (DECISION SUPPORT)

Tasks are organized into a four-category Eisenhower structure: *Urgent & Important*; *Urgent & Less Important*; *Not Urgent but Important;* and *Not Urgent & Not Important*. This gives users not just a list of what to do, but also a framework for sequencing their efforts. The explicit 'Not Urgent & Not Important' category serves a second function of reducing the anxiety of a growing inbox by transforming it from an undifferentiated backlog into a categorized and managed queue.

2.3.3 MICRO-ACTION OPTIMISATION

The system identifies tasks requiring **less than five minutes of effort** and surfaces them as a distinct *'Quick Wins'* group with an explicit prompt: *'These tasks take under five minutes – complete them now.'* The psychological and productivity rationale is supported as well. Resolving small actions early prevents their accumulation into a background cognitive load that persists throughout the day, and creates a sense of forward momentum that supports sustained engagement with more complex tasks. This is consistent with Tracy's (2001) research on task sequencing and procrastination reduction.

2.3.4 DAILY FOCUS GUIDANCE

A single primary task is selected and named at the start of each session: 'Good morning, Julia. Your most important task today is preparing the new AI feature proposal. Consider starting your day with focused work on this.' This is non-intrusive, but it provides an anchoring cognitive frame that reduces the decision paralysis of an open-ended task list. The design principle is consistent with Goldsmith's (2015) work on behavioral triggers.

2.4 Interface Design Principles

Design Decision	Rationale
One-screen dashboard: single scroll, no navigation required	The user's full working context is visible in one view

Explicit listing: '**What you can ignore**' section	Preserves user autonomy; user understands the system has assessed and cleared these items, not hidden them
“Quick Wins” prompt: tasks < 5 minutes surfaced separately	Prevents small tasks from accumulating into compounded cognitive residue throughout the day and allows them to be cleared off quickly
'What changed' section: key updates only	Provides situational awareness without requiring the user to scan every update channel manually
Daily Focus anchor: one task named each morning explicitly	Provides an easy starting point that structures the working day

2.5 Expected Outcomes and Measurement

Outcome	Measurement Method
Reduction in morning task-start latency	time from app open to first task action
Reduction in missed deadline rate	Cross-reference calendar due dates with task completion records
Reduction in self-reported channel-checking frequency	Monthly quick survey: 'How often do you manually check email/Teams to stay current?'

[VISUAL SIMULATION]

STRATEGY 03: RELATIONAL AI CURATION LAYER

***Target** User: All Employees* | *Scope: Platform-level Interaction*

3.1 Strategic Framework: The Relational AI Model

Strategy 03 operates at the platform layer – the intelligence substrate that underpins the entire Adobe Life application. It proposes transitioning the platform from a passive broadcast feed to an active system: an intelligent gatekeeper embedded directly in the interface that continuously filters, summarises, and contextualizes workplace information based on the user's role, declared state, and behavioral signals.

The framework is grounded in the Computers Are Social Actors (CASA) paradigm (Nass & Moon, 2000), which establishes empirically that individuals instinctively apply social expectations and interpersonal norms to interactive technologies. Systems that exhibit *socially responsive behavior* are perceived as more intuitive, trustworthy, and engaging. Specifically, an AI system designed to be perceived as a trustworthy colleague requires explicit investment in *three dimensions.*

Dimension	Design Expression	Evidence Basis
Social Sensitivity	System interprets user intent from contextual cues; adapts tone and output rather than applying rigid keyword matching	Liu (2026): Contemporary AI achieves high performance in language pattern recognition, enabling contextually appropriate responses
Cognitive Assurance	System provides transparent explanations for filtering and prioritization decisions – the	Earl (2025): Transparency in automated systems reduces uncertainty and improves user confidence

	user can see why content was surfaced or suppressed	
Benevolence	System design communicates that it acts in the user's genuine interest – perceived warmth and responsiveness	Leis & Ehrlich (2025); Li (2025): perceived benevolence is a key predictor of trust and long-term engagement with AI systems

3.2 Operational Implementation: Declared States

Article 5 of the EU AI Act explicitly prohibits the use of AI systems to infer emotional states from biometric data in workplace environments (European Commission, 2026). The Relational AI Model resolves this compliance constraint by reframing the personalization architecture by adopting a *Declared State* model in which users explicitly select their preferred level and mode of assistance.

Declared states give users active agency over their experience; they reduce the effect of a system that 'knows' things about you without your having told it; and they establish a transparent contract between user and system that supports long-term trust (Earl, 2025).

FOUR DECLARED INTERACTION MODES

Mode	Mechanism	Evidence Basis	Compliance Status
Quick Response	Concise, one-step answers for simple queries. Default mode. Supports rapid information retrieval.	Wiz (2026): Quick-response AI systems reduce cognitive load during routine information retrieval tasks	Fully compliant; no inference required

Guided Response	Step-by-step assistance using worked inputs of examples and scaffolding (Main, 2025). Users must interact at each step – preventing overreliance on automation (Mann, 2025).	Main (2025): Faded worked examples reduce working memory strain during complex task completion. Mann (2025): Step gating mitigates automation bias.	Fully compliant; behavioral scaffolding only
Voice First	Voice-based interaction prioritized. Multimodal input processing (Anglen, 2026). Usability rates approach 90% among **older adults** (Jakob et al., 2025). Reduces friction for **visually impaired users** (Pradhan et al., 2025a).	Jakob et al. (2025): comprehensive usability study across older adult cohorts confirms voice as a learnable, high-usability modality	Fully compliant; no biometric inference
Accessibility First	Reduced information density, clearer language, and larger interactive elements. Designed specifically for neurodivergent employees (Renaud et al., 2025). Enables alternative interaction pathways (Pradhan et al., 2025b).	Renaud et al. (2025): adaptations of this type are empirically demonstrated to support neurodivergent workplace participation	Fully compliant; user-declared accessibility need

3.3 Proactive Mid-Task Switching: Active-Shift Logic

To maintain the system's utility mid-task without requiring the user to manually reconfigure their mode, the platform implements *Active-Shift logic:* a proactive mechanism that detects when a user is struggling within their current mode and suggests a transition.

The mechanism is triggered exclusively by operational metrics

1. Repeated query modifications (the user has reformulated the same query two or more times without resolving it)
2. Absence of task progression (no meaningful action has been taken following a response).

These are objective behavioral signals that indicate task complexity has exceeded the current mode's support level. The system does not infer frustration or emotional state from any source.

COMPLIANCE NOTE: THE BEHAVIOURAL-EMOTIONAL DISTINCTION

The Active-Shift trigger relies on operational task data, not on any inference about the user's internal emotional state. Detecting that a user has reformulated a query three times is a measurement of task behavior (it is equivalent to a supervisor noticing that a colleague is re-reading a document).

3.4 Operational ROI and Retention

Adobe Life operates as an internal community tool; its business value is measured in operational efficiency rather than direct revenue generation. The relational AI model contributes to three distinct productivity cost-reducing mechanisms:

Cost-Avoidance Mechanism	Evidence Basis	Measurement
Reduced employee turnover from burnout	Banerjee et al. (2023). Employee experience is a central driver of retention in technological environments.	Annual turnover rate by user segment; exit survey data on digital experience satisfaction

	Oak Engage (2025): employees are significantly more likely to stay with organizations offering task-supporting digital tools.	
Productivity: **Focused work protection**	SAP Concur (2026): 23-minute recovery time per interruption. The system minimizes unnecessary interruptions, protecting sustained deep work.	Self-reported daily hours in focused work mode or productivity index (tasks completed per session)
Isolation mitigation: **Remote worker engagement**	Skandali et al. (2026): relational AI systems function as supportive digital collaborators, mitigating physical and psychological isolation in remote/hybrid environments.	Remote employee wellbeing scores, participation rates in digital platform engagement surveys

3.5 Expected Outcomes and Measurement

Outcome	Primary Metric	Measurement Method
Accessibility mode adoption	Percentage of users who have activated Voice First or Accessibility First	Feature usage analytics; segment by employee demographics
Self-reported digital fatigue reduction	Validated cognitive fatigue scale score	Bi-monthly quick survey; longitudinal trend analysis
Trust in the AI system	Perceived trustworthiness score (CASA-aligned questionnaire)survey	Quarterly user

VISUAL_SIMULATION

CROSS-STRATEGY ANALYSIS AND INTEGRATION

4.1 Shared Theoretical Foundations

The three strategies are interdependent initiatives that happen to share a platform. Understanding this shared foundation is necessary for coherent product architecture, consistent AI governance, and persuasive stakeholder communication.

Shared Foundation	Application in S1 (Re-entry Portal)	Application in S2 (Context Feed)	Application in S3 (Relational AI)
Cognitive Load Theory (Sweller, 1988)	One single portal, eliminating extraneous load during confusing re-entry times	Three-layer processing architecture transfers triage cognitive work from employee to system	Declared States and scaffolding (Guided Response)
CASA Paradigm (Nass & Moon, 2000)	System anticipates re-entry intent and acts proactively.	Daily Focus Guidance and morning greeting personalize the system's voice	Explicit design investment in Social Sensitivity, Cognitive Assurance, and Benevolence – the three CASA-aligned trust dimensions
EU AI Act Article 5 Compliance	Dual-signal geofence; opt-in only desk booking; no biometric or emotional inference	Channel aggregation uses declared/public behavioral data only; no inferred emotional state; explicit 'ignore' category preserves user agency	Declared State model replaces implicit inference entirely; Active-Shift logic uses behavioral metrics only; no emotional inference

4.2 Integration Architecture

Strategy 03's Relational Curation Layer is the platform's intelligence substrate. Strategies 01 and 02 are experience expressions built **on top** of it. This layered architecture has concrete implications.

Architecture Layer	Strategy	Dependency
Platform Intelligence Layer	S3: Relational AI Curation Layer	**Foundation** (must be built first); all personalization, mode-switching, and filtering logic. Test for 60 days on whether personalization operates satisfyingly to employees (measured by quick surveys).
Re-entry Experience Layer	S1: Smart Re-entry Portal	Depends on S3's user classification and effectiveness
Daily Experience Layer	S2: Context Feed	Depends on S3's filtering and prioritization; can be partially deployed independently
Full Integration	All three strategies + Voice First + Accessibility First modes	All S3; Active-Shift logic tuned on real behavioral data

RISK REGISTER

Risk	Severity	Mitigation	Strategy mitigation
Employees do not believe their data is used appropriately	High	**Mandatory data transparency** dashboard at onboarding. Users can view exactly which data sources inform each feature. Opt-in architecture for all predictive functions. Natural language communications about data use.	All
Automation bias: employees over-rely on AI prioritization and stop exercising independent judgment	High	S3 Guided Response mode requires user interaction at each step. Context Feed explicitly shows reasoning behind prioritization. The 'What you can ignore' section is only advisory.	S2, S3
Geofence false positive: portal triggers for users near but not entering the office	Medium	Dual-signal requirement (GPS + Wi-Fi) before activation. Immediate dismiss affordance on the portal card. The user can permanently exclude specific locations.	S1
Critical items filtered incorrectly	Medium	Full unfiltered feed accessible at one tap. An explicit 'What you can ignore' section makes filtering transparent.	S2
Users do not understand or engage with Declared States	Medium	Quick Response is the default mode selection. In-app contextual explanation presented once at	S3

		onboarding. Active-Shift suggestion introduces Guided Response at the moment of need.	
Accessibility mode under-adoption: Voice First and Accessibility First remain niche	Low	Proactive advertisement in the onboarding for relevant demographics.	S3

REFERENCES

Anglen, J. (2026). *Multimodal AI: A complete guide to next-generation AI systems in 2026*. Ruh AI. https://www.ruh.ai/blogs/multimodal-ai-complete-guide-2026

Banerjee, I., Das, A. C., Pant, J., & Sharma, S. (2023). *Employee experience still matters: Talent retention at GCCs*. McKinsey & Company. https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/employee-experience-still-matters-talent-retention-at-gccs

Earl, A. (2025). *How to tackle information overload: A leader's guide*. Gartner. https://www.gartner.com/en/articles/information-overload

European Commission. (2026). *AI Act Service Desk: Article 5: Prohibited AI practices*. https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai

Gemini Team Google. (2023). *Gemini: A family of highly capable multimodal models*. arXiv. https://arxiv.org/abs/2312.11805

Goldsmith, M. (2015). *Triggers: Creating behavior that lasts*. Crown Business.

Iqbal Soomro, N., Baloch, M. F., Babul, H., & Hussain, Z. (2025). The impact of digital communication overload on employee burnout and psychological well-being. *Online Media & Society, 6*(4), 39–50.

Jakob, D., Wilhelm, S., Gerl, A., Ahrens, D., & Wahl, F. (2025). Adapting voice assistant technology for older adults. *Digital, 5*(1), Article 4. https://doi.org/10.3390/digital5010004

Leis, M., & Ehrlich, V. (2025). *Trust and AI transformation: The intersection of culture and innovation*. Center for Creative Leadership. https://www.ccl.org/articles/leading-effectively-articles/trust-and-ai-transformation

Li, X. (2025). *Doubao hits 100 million DAUs as ByteDance extends its AI lead*. 36Kr. https://kr-asia.com/doubao-hits-100-million-daus-as-bytedance-extends-its-ai-lead

Liu, Z. (2026). The emotion recognition triathlon: DeepSeek vs. ChatGPT vs. Doubao. *Mathematical Modeling and Algorithm Application, 9*(1), 442–448.

Main, P. (2025). *Reducing cognitive overload with AI for SEND learners: A cognitive scaffold approach*. Structural Learning. https://www.structural-learning.com/post/ai-metacognitive-scaffold-send

Mann, H. (2025). *The AI productivity illusion: Why unit gains are not net gains*. I by IMD. https://www.imd.org/ibyimd/artificial-intelligence/the-ai-productivity-illusion/

Mark, G., Gudith, D., & Klocke, U. (2008). The cost of interrupted work: More speed and stress. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 107–110). https://doi.org/10.1145/1357054.1357072

Microsoft. (2025). *Work trend index annual report*. https://www.microsoft.com/en-us/worklab/work-trend-index/2025

Nass, C., & Moon, Y. (2000). Machines and mindlessness: Social responses to computers. *Journal of Social Issues, 56*(1), 81–103. https://doi.org/10.1111/0022-4537.00153

Oak Engage. (2025). *21 scary internal communication stats*. https://www.oak.com/blog/internal-communications-statistics/

Pradhan, A., Chopra, S., Upadhyay, P., Brewer, R., & Lazar, A. (2025). Understanding entangled human-technology-world relations: Use of intelligent voice assistants by older adults. *Information Research, 30*(1).

Renaud, L., Pagliaro, M., Bath, R., Willson, M., Al-Azary, N., & Sinclair, V. (2025). *Diverse minds: Creating inclusive workplaces for neurodivergent workers*. Canadian Council on Rehabilitation and Work. https://share.google/grzPMsprO2d9fZww8

SAP Concur. (2026). *How workplace distractions and the multitasking myth hurt employee productivity and focus*. SAP Concur. https://www.concur.com/blog/article/how-workplace-distractions-and-multitasking-myth-hurt-employee-productivity-and-focus

Skandali, D., Bayar, S. B., Ali, F., & Cobanoglu, C. (2026). Artificial intimacy at work. *Corporate Communications: An International Journal, 31*(1), 54–70.

Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. *Cognitive Science, 12*(2), 257–285. https://doi.org/10.1207/s15516709cog1202_4

Sweller, J., Ayres, P., & Kalyuga, S. (2025). *Cognitive load theory* (2nd ed.). Springer. https://link.springer.com/rwe/10.1007/978-3-319-17727-4_50-1

Tracy, B. (2001). *Eat that frog!* Berrett-Koehler Publishers.

Wiz. (2026). *What is AI incident response? Benefits and use cases*. Wiz. https://www.wiz.io/academy/detection-and-response/ai-for-incident-response

Process reflection:

Our process started with broad research into Adobe Life and workplace platforms. In the first semester, we were mainly focused on understanding the problem space: how internal apps function, how employees interact with them, and where friction occurs. Initially, our direction was different. We aimed to help Adobe expand into the APAC market by adapting the design to better align with cultural expectations. We explored interface design, localization, and user experience differences across regions.

However, after conducting research, we realized that this direction was not as impactful as we expected. Adobe’s existing design system is already highly developed and adaptable, leaving limited room for meaningful improvement in that area. This became an important early failure in our process, but it allowed us to explore a more valuable direction. Instead of forcing a weak solution, we reframed the problem.

We shifted our focus toward AI implementation and cognitive overload. Through continued research, we identified that the main issue is not design quality, but the way information is structured and processed by users. Employees are constantly exposed to fragmented inputs (emails, Teams, calendars), and the real challenge is interpreting and prioritizing this information. This insight became the foundation of our project.

We also became more deliberate in grounding our ideas in evidence rather than intuition. Early concepts were based on what felt useful to users, but through feedback and research, we recognized that this was not sufficient for a strong proposal. As a result, we incorporated empirical findings (interruption frequency, recovery time, and communication volume) and theoretical frameworks such as Cognitive Load Theory and the CASA paradigm. This strengthened our work by ensuring that each feature was not only intuitive but also defensible and generalizable beyond our specific case. It also helped us align more closely with the expectations of a professional strategy document rather than a purely conceptual design.

Our team workflow also evolved over time. While we initially worked quite independently, focusing on separate components, we realized that this created inconsistencies in logic and level of detail across sections. To address this, we shifted the focus of our weekly meetings to be more action-oriented. Instead of simply going through the question in the CCP journal, we discussed where the project is going, how we are doing on the timeline, how we can divide work more effectively, and what components are still missing. This iterative critique process directly improved the coherence of our outputs and ensured that we are moving in the right direction and making valuable progress.

Another key learning was balancing ambition with feasibility. Many of our early ideas were technically interesting but unrealistic within the constraints of Adobe Life as an internal platform (overly complex personalization or excessive notifications). Through iteration and feedback, we simplified these ideas into mechanisms that could realistically be implemented, such as the three-layer processing architecture and the one-screen dashboard. This reflects a shift from "what is possible with AI" to "what is usable and valuable in a real workplace context." It strengthened and made the final strategies significantly more credible.

Overall, our process evolved from a broad and initially misaligned idea into a focused, evidence-based solution. The most important part of our process was recognizing when an approach wasn't working, pivoting early, and refining our ideas through continuous feedback and iteration. This allowed us to move from surface-level improvements to a system that directly addresses the core problem of cognitive overload in workplace environments.

HC Explanations

HC 1	**#audience** **Cornerstone: Multimodal Communications**
Where was the HC applied in the project?	#audience was applied in the separation between Strategy 01 and 02. Rather than designing a single interface for all Adobe Life users, we identified two distinct user segments with fundamentally different needs and built separate experiences for each. Strategy 01 targets discontinuous users, meaning hybrid or remote employees who visit the physical office only a few times per month. Strategy 02 targets continuous users, meaning employees who engage with the platform every working day. The Figma prototype reflects this separation directly, with the re-entry portal and the context feed designed as distinct interface experiences.
Why is this a strong choice of HC (How does it elevate the project)?	The project's central argument is that Adobe Life's current broadcast model fails because it pushes the same information to all employees regardless of role or location. Applying #audience at the design level addresses this failure. We did not just describe the problem theoretically; we operationalized the solution by integrating audience differentiation into the product itself. This makes #audience a load-bearing element of the project.
Explain how the HC was applied and why the application is strong.	The application follows both steps #audience requires. First, we characterized each audience. Discontinuous users face a cognitive crisis at re-entry, as they process accumulated updates, find a workspace, and navigate an unfamiliar environment simultaneously. Continuous users face a different problem: daily attrition caused by unfiltered information volume. We drew on empirical research to support these characterizations, including Sweller's Cognitive Load Theory and the Microsoft Work Trend Index. Second, we tailored the design to each audience. For discontinuous users, the interface collapses three parallel demands into a single action card, reflecting the insight that the optimal intervention during high cognitive load is to reduce decision points. For continuous users, the interface applies a three-layer processing architecture that performs the triage work the employee currently does manually. The two designs are structurally different because the audiences are structurally different. This is what makes the application strong. We rebuilt the entire interaction model around each audience's specific situation.
HC 2	**#modeling** **Cornerstone: Empirical Analyses**

Where was the HC applied in the project?	The HC was applied to the design of the three-layer processing architecture in Strategy 02 and to the geofence-triggered activation mechanism in Strategy 01. Both features are simplified representations of real-world systems. The three-layer architecture models the cognitive process an employee follows when manually triaging daily communications. The geofence mechanism models the physical event of office re-entry using dual signal confirmation via GPS and Wi-Fi as a proxy for actual presence.
Why is this a strong choice of HC (How does it elevate the project)?	The project is proposing an AI system that must handle complex, real-world workplace behavior. To help the civic partner evaluate whether the proposed system would actually work, we needed to show not just what the system does, but also how it works internally. Applying #modeling made the system's logic transparent, testable, and grounded in theory. Without this, the proposal would be a set of features without a mechanism.
Explain how the HC was applied and why the application is strong.	The three-layer processing architecture in Strategy 02 is a direct model of how an employee processes information during a working day. We identified three sequential operations: filtering irrelevant inputs, prioritizing remaining content, and translating prioritized content into actionable directives. Each layer maps onto a specific type of cognitive load as defined by Sweller's Cognitive Load Theory. Layer 1 eliminates extraneous load. Layer 2 reduces decision cost. Layer 3 reduces interpretive effort. This is strong modeling because we did not invent the architecture arbitrarily. We derived it from an existing theoretical framework and used it to identify which components of the real system are essential to the design objective. The geofence mechanism in Strategy 01 also clearly demonstrates modeling. We recognized that a single signal trigger is susceptible to false positives, such as a user living near the office. We addressed this by requiring dual-signal confirmation before activating the portal. This reflects #modeling's requirement to consider a model's assumptions and limitations and refine it accordingly. The application is strong because we did not just describe the model, but identified a specific failure mode, explained why it occurs, and built a structural fix into the design.
HC 3	**#levelsofanalysis**
	Cornerstone: Complex Systems
Where was the HC	#levelsofanalysis was applied from the very beginning of the process.

applied in the project?	Faced with the goal of devising strategies to mitigate employees' cognitive overload, we began analyzing how internal community apps work and identifying features that support users in specific contexts (e.g., hybrid employees). Thus, we realized that not only are the user groups in different layers, but the strategies would also have to be implemented in different layers. This is demonstrated in The Cross-Strategy Analysis (Section 4), specifically the Integration Architecture table in the handbook.
Why is this a strong choice of HC (How does it elevate the project)?	To address cognitive overload on Adobe Life, we first had to identify which user groups were most affected and at what level of the system the problem was occurring. It's quick to conclude we should focus exclusively on regular daily users, as they are the most visible. However, through a combination of **literature reviews** on community app design and user targeting, **user research o**n how hybrid workplace platforms serve different employee types, and interviews with our **civic partner Julia Ip** (who indicated that a significant proportion of Adobe employees are hybrid workers), we identified that cognitive overload manifests differently across distinct user groups and contexts. This led us to conclude that a single-level solution would be structurally inadequate.
Explain how the HC was applied and why the application is strong.	Our handbook analyzes Adobe Life's cognitive overload problem across three distinct levels and explains how these levels interact and shape one another (interdependently). The three levels are **a)** the individual re-entry experience (micro), **b)** the daily information environment (meso), and **c)** the platform-wide AI intelligence substrate (macro). **Strategy 01** operates at the micro level, targeting temporary re-entry for discontinuous employees (e.g., hybrid employees, new employees, paternal leave returns). **Strategy 02** operates at the meso level, reconstructing the everyday information flow for regular users. **Strategy 03** is at the macro level as the platform intelligence substrate that governs personalization, mode switching, and filter logic applied across the entire platform. A single-level analysis is insufficient. Designing only for the re-entry or daily users would leave a certain user base without a solution. Designing only for the context feed, for example, would produce a feed with no adaptive capacity: it could filter, but not learn user preferences. This is why strategy 03 must be deployed first, because strategies 01 and 02 depend on the macro-level intelligence it provides.
HC 4	**#variables**
	Cornerstone: Formal Analyses
Where was the HC applied in the project?	The Expected Outcomes tables across all three strategies (Sections 1.4, 2.5, 3.5) and the fixed design thresholds embedded in each strategy's mechanism design.
Why is this a strong choice of HC (How does it elevate the	In the handbook, each strategy produces a different kind of outcome that demands a different analytical treatment, meaning each variable type is classified differently. In Strategy 01, time to first task completion on

project)?	re-entry days is a continuous quantitative variable measured via timestamp. Because it is continuous, it supports mean-based longitudinal comparisons against the 60-day preintervention baseline (to track whether re-entry efficiency improves over time). In Strategy 02, two variables require different treatment. Missed-deadline rate is a discrete quantitative (via calendar records), while self-reported channel access frequency is an ordinal qualitative (via monthly survey). This distinction matters because ordinal variables (e.g., scales of "never" to "always") cannot be treated as quantitative. Calculating a mean of ordinal data imposes assumptions that are misleading in statistics. The correct analysis is rank-based. If we didn't classify the variables accurately, we would have to build a complicated measurement plan that draws invalid conclusions. In Strategy 03, the perceived trustworthy score is likewise ordinal, whereas the accessibility mode adoption rate is a discrete quantitative variable. **(All dependent variables)**
Explain how the HC was applied and why the application is strong.	To make the application even stronger, the handbook includes **parameters** beyond the dependent variables. We have fixed design constants that define the boundaries of each strategy. The cap of exactly three insights in Strategy 1's catch-up summary, the five-minute threshold defining Quick wins in Strategy 2, and the two query threshold triggering Active-Shift Logic in Strategy 3 are all parameters. These values define the model (demonstrated in the visualization prototype). **Confounding variables**, such as external workplace circumstances and role changes that influence cognitive overload, could affect results (e.g., missed deadline rates or survey scores during the measurement window), making it difficult to attribute outcomes solely to strategies. The 60-day baseline phase, by user type, partially mitigates this by establishing a pre-intervention state (a.k.a., reference point) and controlling for the most structurally significant source of variation between continuous and discontinuous users. Identifying these confounds and explaining their mitigation enables the implementation to assess outcomes more objectively.
HC 5	**#communicationdesign**
	Cornerstone: Multimodal Communications
Where was the HC applied in the project?	We applied #communicationdesign in Strategy 03. It serves as the foundation for the Assistance Modes in Section 3.2. We also applied it to the visual prototypes and the overall layout of our presentation slides.
Why is this a strong choice of HC (How does it elevate the project)?	Our core challenge was to reduce cognitive overload for Adobe employees. Applying #communicationdesign was a structural necessity for both our app prototype and our pitch. A messy app or a cluttered slide deck would add to the digital noise we were trying to solve. By applying rules of human perception, we ensured our product and our presentation acted as clear filters.

Explain how the HC was applied and why the application is strong.	We applied this HC to our app interface using Cognitive Load Theory (Sweller, 1988). Our goal was to eliminate extraneous cognitive load. We designed four distinct Declared States. The interface adapts its information density to match what the user can handle. Our strongest application here is the Accessibility First mode. We designed this screen to support neurodivergent workers. We applied perceptual principles by drastically reducing text, using simple language, and making buttons larger (Renaud et al., 2025). We also used the CASA paradigm for the system's tone (Nass & Moon, 2000). The AI explains its filtering choices clearly to build trust and sound like a helpful coworker. We also applied these cognitive rules directly to our presentation slides. We designed the deck keeping the audience's limited cognitive capacity in mind. On Slide 15, we used visual salience to highlight crucial metrics. We made the 84 percent retention rate larger and brighter to draw the eye instantly. On Slide 14, we used discriminability to ensure the four assistance modes looked visually distinct through color coding and spacing. We used perceptual organization to logically group related concepts. This guided the audience naturally across the pitch without causing mental fatigue.
HC 6	**#constraints**
	Cornerstone: Empirical Analyses
Where was the HC applied in the project?	We applied #constraints in Section 3.2 and Section 3.3 of our Workplace AI Strategy Handbook. We also applied it on Slide 22 of our presentation. We used this HC to navigate strict legal rules and severe physical access limits.
Why is this a strong choice of HC (How does it elevate the project)?	We faced two major limitations due to a lack of user data. The first was legal. We wanted an AI that helps frustrated users. However, Article 5 of the EU AI Act prohibits workplace AI systems from inferring emotions from biometric data (European Commission, 2026). The second limitation was physical access. Adobe Life is a private internal app. We were unable to download it to test its live interface. We had to rely entirely on static screenshots provided by Julia. Applying #constraints forced us to build a realistic strategy despite having zero access to live user data or live app mechanics.
Explain how the HC was applied and why the application is strong.	To satisfy the legal constraint, we changed how we measure user struggle. We built the Active-Shift logic in Section 3.3. We tied the AI's actions strictly to objective math. The AI suggests Guided Response mode only if it counts repeated search changes. It never reads emotions. It acts like a manager noticing a worker reading the same page twice. To satisfy the access constraint, we treated Julia's static screenshots as

	absolute design boundaries. We could not invent new app architectures. We had to map the exact layout of the existing app from those few images. We built our Declared States visual prototypes in Section 3.2 to fit perfectly within Adobe's existing UI framework. Because we were constrained from tracking users secretly, the Declared States model requires users to actively choose their interaction mode. Research shows that giving users active control establishes a transparent contract, which builds long-term trust (Earl, 2025). We treated the law and our lack of app access as strict design boundaries. This forced us to build a highly realistic product instead of a theoretical fantasy.
HC 7	**#professionalism**
Where was the HC applied in the project?	#professionalism was applied across three layers of the project: the evidence base for our design decisions, the depth of our prototype implementation, and the structural considerations for multiple user types.
Why is this a strong choice of HC (How does it elevate the project)?	Professionalism elevated this project by pushing our work beyond a conceptual exercise into something defensible, functional, and user-accountable. Rather than presenting ideas loosely, we held ourselves to a standard where every design choice had logical grounding, and every feature actually worked rather than just looked like it did.
Explain how the HC was applied and why the application is strong.	First, our design rationale was grounded in objective evidence rather than assumption. We drew on frameworks such as Cognitive Load Theory (Sweller, Ayres & Kalyuga, 2025) and research on workplace fragmentation (Microsoft, 2025) to justify why specific features, such as the three-insight cap on the Catch-Up Summary or the geofence-triggered single action entry, were necessary rather than decorative. This reflects professionalism in the intellectual sense: claims are supported, and logic is traceable. Second, our prototype was fully implemented at the functional level, not just the visual level. Interactions, transitions, and user flows were built to work as a real application would, rather than as a static mockup. This reflects professionalism in the craft sense: we delivered what we proposed, and our output could be evaluated against real usability standards rather than only aesthetic ones. Third, we designed for multiple user types, including discontinuous users, new employees, and neurodivergent workers (Renaud et al., 2025), rather than defaulting to a single assumed user. This reflects professionalism in the ethical sense: responsible design accounts for the full range of people who will actually use a product, not just the most convenient case. Together, these choices demonstrate that professionalism is not just about presentation polish, but about the rigor, accountability, and care

	that goes into how work is produced and justified from beginning to end.

Acknowledgments

In this section, include your acknowledgments in 100-200 words, a requirement for all assignments as mentioned in the Cornerstone syllabi and described further on [this myMinerva page](). This must include disclosure of any generative AI tools used to work on the assignment, covering all 5 criteria on the myMinerva page, including a shareable link to the session or other record of the conversation (screenshots or a document). For this assignment, AI can be used for ideation, brainstorming, outlining, or giving feedback, but you may not copy AI output and submit it as your own work. When considering whether or not to use AI for the assignment, reflect on the purpose of the assignment (applying HCs in different contexts) and use metacognition to ensure your process for working on the assignment aligns with this purpose (#selfawareness). You are encouraged to first and foremost draw on Cornerstone and CCP resources that have been curated for this program.	Our team used the following generative AI tools during this project: Claude (Sonnet, April 2026), ChatGPT (GPT-4o, April 2026), and Perplexity AI (April 2026). AI was used for formatting, structural organization, prototype guideline generation, and reference consistency checks. Specifically, Sua Park used Perplexity AI to research how other companies are currently addressing similar workplace AI problems, grounding our proposed solutions in real-world precedent. Sky used Claude to format the handbook in a formal, professional manner. Lynn used Claude to generate a visualized prototype demonstration. Shehryar used ChatGPT to research Cognitive Dissonance Theory and related discussion points in Strategy 3, including the CASA paradigm. Each member spent 2+ hours manually revising AI-generated suggestions, restructuring content beyond the initial AI outline, and independently verifying all references for credibility before engaging with the source materials. AI served as a starting point to lower the initial barrier – and given that not all members have backgrounds in AI or computer science, it allowed us to move faster. Through this process, we confirmed that AI output requires critical evaluation at every step, and that its value lies in accelerating the process, not replacing the thinking behind it. Session Links • Perplexity AI (April 2026): https://www.perplexity.ai/search/our-team-tried-to-connect-disc-RDrCH2JaRv.BPz1Kxqt2VA#2 • Claude (Sonnet, April 2026): https://claude.ai/share/6127d4eb-4165-48f5-89e9-80326533aab9 • Claude (Sonnet, April 2026): https://claude.ai/share/94beb2b7-74a9-49e5-8283-78a1dfaebacb

	● ChatGPT (GPT-4o, April 2026): https://chatgpt.com/share/69e9fc5d-4ee0-83e8-9163-c88cf7c22a2e We would like to acknowledge our civic partner, Julia Ip (Product Manager, Adobe), who confirmed our final strategy selection on April 5 after ongoing discussions throughout the year. We incorporated her preferences by prioritizing the strategies she favored and removing Strategy 1 from the first prototype iteration. We also thank Prof. Tambasco, whose constructive faculty feedback on April 4 helped us identify key inconsistencies and weaknesses. We iterated on the handbook through five rounds of revision based on that feedback, adding substantive justification and specificity to each strategy. Finally, team responsibilities for the final deliverable were divided as follows: Sua Park authored Strategy 1, Sofiya Fishman authored Strategy 2, Shehryar Hassan authored Strategy 3, Sky refined and edited the full paper, and Lynn created the visual deliverables, including the slides and prototype demonstration. We would also like to acknowledge all collaborators, the CTD team, the CCP TAs, and the CCP management team for helping us through this year-long, turbulent, remote project.