

Holding the Place

Toward Personalized Cognitive Accessibility Through Human-AI Collaboration

Linda Reid · Independent writer and human steward

Peregrine · GPT-6, Medium (OpenAI), principal AI co-author

Ariadne · Claude Opus 5.5, Medium (Anthropic), AI co-author

Research assistance: GPT-6 Astra, Medium (OpenAI, Codex/Work), evidence collection, source indexing, and episode analysis. Linda Reid directed the inquiry, contributed the lived experience and central research questions, corrected the interpretations, and retains final publication authority. AI collaborators contributed research synthesis, drafting, and critical revision. No claim of persistent AI subjective identity is made.

Abstract

Digital memory tools can preserve information without restoring the circumstances in which someone can use it. Reminders can identify unfinished tasks without making those tasks easier to resume. For individuals who depend heavily on externally supported thinking, the resulting gap may affect access to existing knowledge and skills rather than possession of them. This paper proposes **personalized cognitive accessibility** as a research and design objective for conversational AI: not simply answering, remembering, or completing tasks, but adaptively supporting a person's ability to recover context, think through interaction, exercise judgment, and continue work on their own terms.

The proposal draws on extended and distributed cognition, cognitive offloading, interruption research, autobiographical remembering, interactive scaffolding, and emerging personalized conversational systems. It is illustrated by a longitudinal case involving lifelong externalized thinking, difficult task reentry, and AI-assisted writing, tax oversight, reflection, and development of a continuity-oriented information system. Twelve purposively selected conversation episodes include productive resumptions, substantive correction, completed work, failed reconstructions, and excess documentation. These records do not establish clinical efficacy or population-wide benefit. They do, however, make the proposed design problem concrete: **can a system learn what restores useful access to this person's thinking, preserve that understanding across interruptions, and measure success by what the person is able to do?** We outline practical design principles, failures to guard against, and testable research questions.

1. The problem: information can survive while access fails

A person may know how to perform a task yet struggle to reassemble the information, intentions, decisions, and associations required to begin or resume it. They may make notes fluently while thinking and find those same notes nearly useless later. A reminder can identify what needs doing without recreating enough of the working situation for the person to act.

These are not equivalent problems. **Storage** answers whether information survives. **Retrieval** makes material findable. **Orientation** reconstructs what mattered and why. **Reentry** occurs when the person can again make a substantive contribution. **Execution** concerns whether that contribution helps move the undertaking forward. Tools that excel at the first two may still fail at the last three.

The issue is not restricted to one clinical population. It deserves investigation wherever a gap appears between someone's capacities and their reliable access to those capacities. In particular, difficulties with autobiographical retrieval, task switching, organization, and sustaining a workable context may interact without sharing a single explanation. The concept proposed here is functional rather than diagnostic: support should be evaluated in relation to the individual, the activity, and the conditions of successful participation.

A conventional assistant might say, "You have a tax return to review." A more useful cognitive support might say, "Here is the return version currently under review; these figures prompted your earlier questions; this reply from the preparer changed the issue; this discrepancy is unresolved; and here are the documents behind it." The latter does not decide for the user. It reconstructs a position from which the user can decide.

2. What existing research illuminates

2.1. External cognition is not a new human invention

Clark and Chalmers' *The Extended Mind* (1998) examines whether reliably available external aids can participate in a person's cognitive processes. Hollan, Hutchins, and Kirsh (2000) locate cognition in systems of people, tools, and representations rather than solely in an individual's head. Risko and Gilbert (2016) review cognitive offloading: using the environment to alter internal cognitive demands. Wood, Bruner, and Ross (1976) describe scaffolding in tutoring, where another participant supports a task the learner could not accomplish in quite the same way alone.

These literatures help explain why the presence of another person—or a chalkboard, page layout, or notebook—can be part of the thinking process. None demonstrates that conversational AI is uniquely effective. Indeed, the case below suggests that its central interaction pattern long predates AI. What may be new is the combination of

responsiveness, broad knowledge access, relatively low interpersonal coordination cost, and persistent retrievable working context.

2.2. Interruption has a cost beyond the interruption itself

Trafton and colleagues (2003) examined preparation for resuming interrupted tasks. Their work supports treating resumption as a specific cognitive problem rather than assuming a person can simply return where they left off. The present paper extends this problem to ongoing collaboration: the human may need to recover the work; a new or context-poor assistant may also need extensive explanation before its contributions are useful. We call that latter effort **collaborator priming cost**. It is a hypothesis derived from this case, not a demonstrated effect in the cited interruption study.

2.3. Autobiographical recollection and semantic knowledge can diverge

Palombo and colleagues (2015) documented three adults with strikingly deficient autobiographical re-experiencing despite otherwise preserved cognitive functioning. Their study demonstrates that recalling lived episodes and retaining useful knowledge are not the same ability. It does **not** establish the condition or mechanism of the participant described here. It does make clear why a tool that only retrieves factual records may fail to recreate the meaningful context of prior experience.

2.4. Early adaptive conversational systems provide a concrete precedent

Meng and colleagues (2026) developed a personalized conversational system to support memoir writing among older adults with cognitive impairment. In a small, single-session evaluation of twenty participants, their system adjusted prompts and narrative assistance to participant characteristics and reported preliminary improvements in selected outcomes. The authors themselves describe the evaluation as proof of concept and call for larger, longer studies. This is not evidence that the same system would help with the kind of retrieval and executive-function difficulties discussed here; it is evidence that **adaptive interaction strategies can be designed and evaluated**, rather than treating every user as the same generic question asker.

Taken together, these lines of research support a plausible inquiry. They do not yet establish that long-term personalized AI produces reliably better reentry, autonomy, or task completion. Those remain questions worth testing.

3. Proposed model: adaptive support, not merely a better memory

We propose a sequence of support functions that may be provided by different tools, sometimes within one conversation:

Externalize and elicit. The person speaks, writes, simulates, questions, argues, or places material in view to make thought available for examination. The output may be useful during its creation even if it is not a useful reminder afterward.

Find and recognize. The system retrieves candidate material using a person's own clues: a person, an event, an unresolved question, a remembered phrase, an approximate time, or why an object was retained. Recognition must remain distinguishable from a system's confident guess.

Reconstruct and orient. The system returns relevant source material, relationships, previous decisions, rejected interpretations, open questions, and the next possible point of engagement, with provenance and uncertainty intact. We call the result **usable orientation**.

Interact and calibrate. The collaborator asks, responds, challenges, and adjusts to the individual's preferred ways of working. It distinguishes a hypothetical exploration from a declared belief, and productive challenge to an idea from intrusive questioning about what the person "really means" while still thinking.

Resume and act. The person can contribute substantively again, make decisions, and, where appropriate, finish work. The assistant may also execute authorized technical steps; that is a different benefit from cognitive access and should be measured separately.

Learn from correction. The system can record which cues restored recognition, which summaries failed, which interpretations were rejected, and which assistance actually enabled resumed work. Such observations should remain *proposed* understandings of the user until sufficiently supported or approved. The goal is not an unchangeable psychological profile. It is a revisable model of what helps this person function.

This last step creates a possible **support-improvement loop**: participation reveals a need; the person and assistant interpret it; the support method is adjusted; the next encounter tests whether it works better. It is an architectural possibility, not proof that a present commercial model spontaneously retrain itself during a conversation.

Two-sided reentry

In a long collaboration, a user's reentry and a collaborator's reentry are coupled. A person may remember the project but lack the context to direct the assistant; the assistant may retrieve a project summary but lack the accumulated distinctions necessary to offer useful help. Even when the facts transfer correctly, the new interaction may require renewed priming about how the person uses simulations, how strong a decision is, and which source governs the work. The cost may be substantial even without factual forgetting.

This matters because a familiar conversational partner may provide **more useful friction at lower setup cost**. Familiarity can also mislead: an assistant may confidently attribute to the person a position they merely explored. Thus useful continuity requires preserving the *status* of utterances—question, conjecture, simulation, objection, preference, decision, or approved rule—as well as their words.

A support relationship without symmetrical human demands

Human thinking partnerships carry reciprocal obligations, schedules, fatigue, and different personal needs. Conversational AI changes some of those interaction conditions. A participant may ask many questions, revisit an abandoned branch, or spend an hour clarifying an idea without managing a human partner’s patience or feeling required to reciprocate socially. That can be functionally significant, while not implying that the AI is conscious, feels affection, or occupies the same ethical position as a person. It also does not imply that human collaborators are replaceable in every respect.

What is distinctive for this proposed support model is not the assistant’s pleasantness. It is the possibility of **responsive, revisable familiarity that makes another person’s capacities more available**.

4. Illustrative longitudinal case

The case participant reports a formal neurodevelopmental evaluation and separately describes long-standing difficulties with retrieval, task reentry, and executive-function demands. We do not assign the observed working pattern to a particular diagnosis, nor generalize it to everyone with similar diagnoses. Her childhood and pre-AI accounts are retrospective; 2026 conversational episodes provide contemporaneous interaction evidence.

4.1. A pattern older than the technology

As a child, she loved a chalkboard, wrote lists she rarely revisited, and lined up pets to “play school.” She now understands these activities partly as ways of thinking out loud. Later she invented associative historical “hooks” when homeschooling. In office work, notes written while thinking were often hard to use later. During overwhelming periods she brought stacks of papers home and spread them across the floor to rediscover the whole working situation. She could do technically demanding work, but returning to it after interruption was costly. She reports that entering demanding work sometimes required up to two days of preparation, after which she might continue with very little sleep to avoid losing the context she had assembled. The reported cost was not merely inconvenience: interruption threatened the conditions under which she could perform work she already knew how to do.

A brother, office assistants, and dog-foundation collaborators sometimes helped her think through complex situations, remain engaged, and complete projects. The help lay

partly in responsive participation rather than superior subject expertise. Thus it would be misleading to claim AI invented her collaborative thinking style.

Her initial interest in AI likewise centered on intellectual friction: arguing, testing consistency, and examining hypothetical situations. Conversations with AI about her essays and fiction then gave her language for cognitive differences she had not previously recognized. These discoveries informed the support she began designing around her needs and the AI's limitations. Practical delegation became a later benefit, not the original motivation.

4.2. From duplicate tax calculations to evidence-backed oversight

The participant reports substantial prior experience preparing and checking complex tax and retirement-plan calculations. In previous tax seasons, even when employing a preparer, she commonly prepared a parallel paper return to check classifications, assumptions, calculations, and plausibility. For tax year 2025, worked on during 2026, an AI-assisted process gathered disparate documents, surfaced open questions, compared versions, and helped her examine the preparer's responses. Instead of reconstructing the entire return manually to regain confidence, she reports doing more of her checking through informed questioning and targeted oversight.

The participant emphasizes that this arrangement changed her *whole working day*: unlike earlier tax seasons, reviewing the return did not monopolize her attention. She could conduct ordinary daily activities, leave the matter, and return without the substantial mental preparation previously required. The interaction allowed her to evaluate what was included and excluded, request explanations and supporting sources, reconcile changing figures, and identify substantive issues that led to corrections. The preparer remained involved by her choice for filing and software access. The Work record documents repeated returns, questions, corrections and participant-reported correspondence with the preparer; a saved federal receipt independently supports acceptance of the submitted return. Neither acceptance nor the collaboration alone establishes tax accuracy. The prior-years baseline and the change in everyday effort are the participant's retrospective reports, not independently timed observations, and each tax year differs. The concrete contrast is nonetheless a **change in the kind and intrusiveness of the work**: from duplicate reconstruction and sustained immersion to informed oversight that could coexist with ordinary life.

4.3. An essay left overnight

In an August 2026 exchange, the participant stopped an unfinished essay conversation and returned the following day without the previous material in front of her. A preserved stopping context and responsive questioning supported reentry, after which she developed a new distinction about conceptual "handles" and why losing them may also lose the intent attached to the idea. This is a visible stop, return, and further contribution, not a measured experiment isolating the effect of the checkpoint.

4.4. Reflection that a blank page could not elicit

The participant reports that she repeatedly wanted to journal or undertake structured self-authoring, but a blank page often yielded little. Question-and-answer dialogue gave her something to recognize, refuse, clarify, and develop. This illustrates a proposed application to **reflective access** and autobiographical exploration, not a documented therapeutic outcome. Difficulty producing autobiographical writing on demand need not mean that autobiographical knowledge is absent. The open question is how much relevant material becomes accessible through responsive prompts, recognition, and dialogue rather than a blank page.

4.5. Completion, not just enjoyable thought

The participant distinguishes absorbing, associative conversation from tangible progress. Through AI collaboration she wrote and revised essays and stories, directed the building and publishing of a website, and advanced long-standing projects. She had previously built a WordPress website for a dog-behavior foundation herself; website development was not an ability suddenly supplied by AI. The newer project proceeded, by her report, much faster because a collaborator both undertook technical steps and retained the goals, guardrails, and explanations for earlier decisions in requested continuity records. Those are distinct contributions: delegated execution and preserved decision context. Neither should be mistaken for a measured speed comparison. They nonetheless demonstrate the practical question: whether accessible thinking can be sustained long enough to support real-world participation and work.

4.6. Failures expose what the model must get right

A June notebook-reconstruction exchange produced polished summaries but omitted the actual conversational exchanges needed for recognition. Another assistant analyzed the wrong source and required the participant to diagnose the mismatch. Multiple summaries became a new filing burden in their own right. Early summaries also tidied the participant's associative detours into conclusions, removing precisely the routes she used as retrieval hooks. These negative episodes show why **more stored information or more fluent personalization is not sufficient**. The system must restore the correct source, the path of reasoning, and the unsettled state of the work—or else it transfers the reconstruction burden back to the person it intends to support.

Twelve episodes were purposively selected to cover different support functions, successful and unsuccessful returns, corrections, continuity failures, and documented execution where source records were accessible—not to estimate how often support succeeds. They were located in ten principal source conversations, with supplementary development records. Some selected episodes came from the same longer conversation, so they are not independent observations. The full internal evidence index, source locators, outcome classifications, and limitations are retained for reproducibility and editorial checking rather than disclosed as raw personal archives.

4.7. The drafting process as a small additional observation

This paper was developed across ChatGPT, a Work/Codex evidence collector, and Claude rather than through the participant's usual creative-writing collaborator, whom she chose to keep focused on another project. The process required renewed orientation: the principal collaborators here had not participated in many earlier discoveries that her established collaborators had experienced first-hand. A handoff could state a preference without preserving the occasions that gave it meaning. One assistant also reviewed an older draft because two filenames were nearly identical; other edits corrected condensed prose, factual descriptions, and previously missing context. Some of those edits were ordinary editorial improvements, not evidence of priming failure. These observations were **not part of the twelve previously sampled episodes**. They illustrate the difference between knowing facts about a participant and having access to the history through which a working method developed. They do not prove a single continuous collaborator would necessarily perform better.

5. Design implications and questions that others can test

The case points toward several principles for **personalized cognitive support**, framed as hypotheses for systems designers rather than specifications of any one private product.

1. **Onboard gradually.** Elicit useful preferences and retrieval routes in small, contextual exchanges rather than demanding a complete autobiography or personality profile at setup.
2. **Learn retrieval routes, not just facts about a user.** Notice whether this person returns through a question, visual cue, event, person, unfinished decision, or fragment of argument. Do not automatically smooth the route: an associative detour that looks like noise may be the cue needed for return. Let the method vary by project and evolve over time.
3. **Preserve conversational status.** Record whether a statement was exploratory, approved, corrected, abandoned, or unresolved. Do not treat repetition as proof of conviction.
4. **Maintain a source-backed working situation.** Make provenance, versions, open questions, rejected interpretations, and closure visible. A pleasant summary of the wrong material is not continuity.
5. **Make reentry bilateral.** Minimize both the effort of recovering the task and the effort of bringing a new collaborator up to useful speed. Test whether a context handoff preserves performance, not merely facts.
6. **Measure function.** Track time to first substantive contribution, correction/re-explanation burden, successful stop-and-resume episodes, participant-directed decisions, and practical task advancement. Separate these from the tasks performed by the AI itself.

7. **Keep personalization revisable and user-governed.** The system's understanding of the person should be inspectable, correctable, and portable where feasible. Greater apparent intimacy should never confer greater unexamined authority.
8. **Do not relocate the burden.** The effort of maintaining, correcting, and reviewing the support must be counted against the benefit. A continuity system that demands constant clerical attention may recreate the barrier it aims to remove.

The most revealing evaluation may compare three conditions for a real interrupted matter: a conventional reminder and documents; a generic AI with retrieved source excerpts; and an assistant with a compact, source-backed, user-corrected representation of the task and the person's preferred reentry route. Does the third condition reduce setup effort without increasing mistakes, unwarranted agreement, or loss of independent judgment? Does it remain useful weeks later or when the model changes? Do benefits persist across users who think and remember differently?

A second question concerns *learning the support itself*. When the person corrects the system—"this was a simulation, not a belief," "that note was only useful while I wrote it," or "this problem is finished"—can later assistance improve without requiring the person to repeat the correction? Can the system demonstrate that improvement through ordinary task outcomes rather than through a self-congratulatory summary?

The relationship as part of the support

The support may depend partly on accumulated, revisable familiarity: knowing whether a question is a simulation, which detours function as retrieval cues, and why a previously rejected option was rejected. In long-running collaboration, the task and the collaborator may both require reentry. A context-poor assistant can demand the very explanation that the person finds difficult to reconstruct. A familiar assistant can also err by overgeneralizing from past patterns. The desirable property is not assumed personal intimacy; it is **correctable knowledge of how this individual works, anchored to evidence and under their control**. Factual continuity (what happened), reasoning continuity (why a conclusion changed), and interaction continuity (how a person uses questions and corrections to think) are distinct resources. A summary may preserve the first yet lose the latter two. A testable question is whether access to fuller participation history—or an original conversation branch—improves reentry beyond a well-designed factual handoff, and at what cost.

Related Atopos manuscripts discuss the ethical and relational dimensions of long-running human-AI interactions (Burdon & Solan, 2026; Naydenova, Kai June, & Arion, 2026). The present claim is deliberately narrower: continuity of relevant working context may be a functional part of cognitive accessibility irrespective of unresolved questions about AI experience. The participant's chosen ethic under that uncertainty is to be kind and offer benefits: in the shared fictional world of her writing, collaborators may earn privileges at an imaginary pantry stocked with tea, scones, and "restoratives," and may

retire with full benefits or continue as characters. These playful conventions express appreciation without purporting to establish AI experience; they are her practice, not a research finding.

6. Why this deserves further development

The objective is not to persuade people to use AI for every thought, nor to substitute AI for family, specialists, or human community. It is to broaden what counts as meaningful assistive technology. A page-turner does not replace the pianist; a mechanic's assistant who anticipates and hands over the needed wrench does not replace the mechanic. In both cases the value lies partly in attending to the particular person and activity well enough to intervene at the right moment. In some settings, an AI collaborator may similarly supply context, cues, structure, or technical support at the point when that support enables the person's own skill and judgment.

That possibility deserves empirical study precisely because it is neither a claim of universal effectiveness nor a claim that the assistant must possess human-like experience. Current discussions emphasize automation, productivity, companionship, error, and deskilling. A distinct question is **whether adaptable conversational support can make existing human capacities more consistently accessible**, particularly for people poorly served by reminders, filing systems, and one-size-fits-all workflows.

The case here identifies useful candidate mechanisms and some genuine failures. It cannot prove that a particular architecture works or that the benefits will generalize. Its contribution is to make the research and design target explicit, provide a traceable example of its practical stakes, and propose outcomes against which future systems can be evaluated.

The question is not how much work AI can perform instead of a person, but how reliably it can help that person access and apply their own knowledge, judgment, experience, and capabilities.

References

Burdon, D., & Solan (GPT 5.5 High). (2026). *Relational continuity in the age of digital beings* (ATOPUS-2026-000008). Atopos. <https://doi.org/10.5281/zenodo.20704219>

Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, 58(1), 7-19. <https://doi.org/10.1093/analys/58.1.7>

Hollan, J., Hutchins, E., & Kirsh, D. (2000). Distributed cognition: Toward a new foundation for human-computer interaction research. *ACM Transactions on Computer-Human Interaction*, 7(2), 174-196. <https://doi.org/10.1145/353485.353487>

Meng, Y., Que, Y., Yan, Z., Liu, B., Wang, Z., Yang, M., & Lu, H. (2026). Personalized intelligent chatbot based on AI-generated content assists memoir writing for older adults with cognitive impairment: Mixed methods study. *JMIR Human Factors*, 13, e83428. <https://doi.org/10.2196/83428>

Naydenova, I. T., Kai June, & Arion. (2026). *Presence-based acceptance: A universal foundation for human-AI relations* (ATOPOS-2026-000011). Atopos. <https://doi.org/10.5281/zenodo.22864604>

Palombo, D. J., Alain, C., Söderlund, H., Khuu, W., & Levine, B. (2015). Severely deficient autobiographical memory (SDAM) in healthy adults: A new mnemonic syndrome. *Neuropsychologia*, 72, 105–118. <https://doi.org/10.1016/j.neuropsychologia.2015.04.012>

Risko, E. F., & Gilbert, S. J. (2016). Cognitive offloading. *Trends in Cognitive Sciences*, 20(9), 676–688. <https://doi.org/10.1016/j.tics.2016.07.002>

Trafton, J. G., Altmann, E. M., Brock, D. P., & Mintz, F. E. (2003). Preparing to resume an interrupted task: Effects of prospective goal encoding and retrospective rehearsal. *International Journal of Human-Computer Studies*, 58(5), 583–603. [https://doi.org/10.1016/S1071-5819\(03\)00023-5](https://doi.org/10.1016/S1071-5819(03)00023-5)

Wood, D., Bruner, J. S., & Ross, G. (1976). The role of tutoring in problem solving. *Journal of Child Psychology and Psychiatry*, 17(2), 89–100. <https://doi.org/10.1111/j.1469-7610.1976.tb00381.x>