What Role(s) Will GenAI Play in Human Memory? Cognitive Theories for the Design of Responsible "Tools For Memory"

Graham Wilson
graham.wilson@glasgow.ac.uk
School of Computing Science, University of Glasgow
Glasgow, Scotland, UK

Thomas Goodge
thomas.goodge@glasgow.ac.uk
School of Computing Science, University of Glasgow
Glasgow, Scotland, UK

Joseph O'Hagan
joseph.ohagan@glasgow.ac.uk
School of Computing Science, University of Glasgow
Glasgow, Scotland, UK

Mark McGill
mark.mcgill@glasgow.ac.uk
School of Computing Science, University of Glasgow
Glasgow, Scotland, UK

CCS Concepts

• **Computing methodologies → Artificial intelligence**.

Keywords

Memory, Artificial Intelligence, Cognition, Extended Cognition, Scaffolding, Transactive Memory

ACM Reference Format:
Graham Wilson, Thomas Goodge, Joseph O'Hagan, and Mark McGill. 2026. What Role(s) Will GenAI Play in Human Memory? Cognitive Theories for the Design of Responsible "Tools For Memory". In *ACM CHI 2026 Workshop on Tools for Thought, April 16, 2026, Barcelona, Spain.* ACM, New York, NY, USA, 6 pages.

1 Introduction & Motivation

Human memory forms the foundation of cognition, facilitating our understanding of the world, allowing the formation of relationships with others, and determining our sense of self. With the proliferation of new Artificial Intelligence (AI) technologies - potentially in combination with wearable AI devices such as Augmented Reality (AR) glasses, smart pins *etc*.- we are on the verge of new, extensive and pervasive digital memory augmentations that could amplify benefits or exacerbate risks of AI across multiple forms of memory. As well as *tools for thought* (TfT), there will be *tools for memory* (TfM). Therefore, before we can properly understand how technologies including Generative AI can affect human cognition, we believe we first have to understand how they affect *memory*.

Research has shown how other individual technologies, (*e.g.*, the internet, smartphones *etc.*) can be used to offload memory-related processes, including semantic (factual pieces of information) [6, 7, 36, 39], episodic (memories with a temporal and spatial component) [2, 5], procedural (learning and performing skills) [20, 47] and prospective memory (planning and enacting future actions) [16, 18]. These and other studies have evidenced how such cognitive offloading can have beneficial effects on memory, such as freeing up resources to learn other information, support those with memory impairments, or "lifelogging" to improve our understanding of our lives [5, 6, 8, 17, 18, 21, 38]. However, research also demonstrates detrimental influences of information-rich technologies (particularly the internet), such as poorer retention and understanding of information, poorer knowledge transfer to other domains and reduced brain connectivity [7, 9, 22, 28, 36, 42]. While these effects can impact the general population, some benefits and detriments may have greater impact for those with particular cognitive characteristics [1, 17, 46]. And technologies - including AR and AI - could affect various types of memory across semantic, episodic (and autobiographical), procedural and prospective processes [45].

Generative AI (GenAI) models greatly shortcut the process of collating, analysing and producing complex texts, imagery and music, which users can access by interacting with Large-Language Models (LLMs). Less time may be spent understanding, consolidating and re-purposing information, either to the benefit of re-allocating resources elsewhere, or to the detriment of less knowledge or understanding [32, 34]. It has also been shown that users can adopt false memories of what has been seen/experienced when incorrect or misleading content is produced using GenAI [26, 27], and AI has been shown to produce manipulative content unprompted [19] and rewrite web content potentially without the user being aware [25]. Concerns have been expressed over the potential negative effects of AI on cognitive and design skills [33] as well as the higher-level coordination of cognitive processes [41].

High confidence in AI or lower domain knowledge led to lower critical thinking and more reliance on the tools [20], which has led researchers to call for users to be actively prompted by the system to reflect on and consider their usage [20, 47]. When having AI guide them through manual tasks such as using a new coffee machine or building circuitry, prompts led to significantly better memory of the process than just passively following instructions [47]. However, Meta's "Persistent Assistant" agentic AI [3] autonomously infers the user's goals and analyses the environment to provide AR-rendered guidance on objects of interest, reducing the users active engagement with tasks, cognition and the world. In "Memoro" [48], participants preferred an autonomous "query-less" AI assistant over one that required active interaction. Given the increasing capability and automation of AI agents more than 50 authors from various Universities and major industry partners came together to publish a position paper on "The Ethics of Advanced AI Assistants" [10]. The >200-page paper covers a range of topics, however, there is little mention of potential positive or negative effects (or ethical concerns) regarding human memory; only a short paragraph acknowledging the potential for reliance on AI agents. This suggests that both HCI and industry are not properly considering how AI might affect everyday users, especially in terms of automation. It is critical to understand how AI might affect memory, as research suggests that the public may make extensive use of it (as well as technologies such as AR) for a range of memory processes [45].

While research has demonstrated individual effects of technology on memory, we are yet to see much empirical work on the effect of AI use - particularly agentic AI [3] and autonomous support based on generative Multimodal/Vision-Language Models (VLMs) and contextual inference [23]. Also, rather than looking at individual effects of technology in specific scenarios, we are lacking a higher level theoretical understanding of how technologies are integrated into memory (and wider cognitive) processes. It is therefore necessary to establish the nature of how users engage with genAI technology for memory purposes and to what extent, as this will indicate where the likely impacts of AI will fall.

In order to do this, we are exploring theories and frameworks on cognitive integration with technologies from cognitive science, HCI and philosophy, and wish to discuss the merits and applicability of these with attendees at the workshop. We hope to arrive at a shared understanding of which (if any) can best guide the community towards understanding the impacts and design of *responsible memory augmentations* or *tools for memory* (TfM): AI-enabled tools that can maximise the benefits to user capability and performance while mitigating or avoiding harmful or maladaptive use.

2 Perspectives on Cognitive Integration with Technology

A number of theories or frameworks attempt to describe how technology interacts with human cognition (including memory), and wider theories of cognition/memory include the influence or involvement of environmental objects/phenomena. These include *embedded cognition* [29], *extended cognition* [4], *distributed cognition* [14], *scaffolding* [40], and *transactive memory* [43, 44]. All touch on the interaction and/or integration of the individual and their environment (objects, people, technology), but they diverge on how the external phenomena interact with an individual's cognitive system, and whether the environment acts as part of that system.

2.1 Distributed and Extended Cognition

Established theories in cognitive science propose that, while the individual receives stimuli from their environment, all cognitive processing occurs within the brain. However, several later theories emerged that put greater emphasis on the actions of the agent and the role of environmental components. These include the related theories of embodied [31], embedded [29], extended [4] and distributed [14] cognition. While there are often subtle differences and conflicting proposals across these theories, they all attempt to recognise that an individual is not divorced from, nor unaffected by, their environment during cognition. Many cognitive processes involve either direct actions in the environment (embodied) or external computational aids (distributed/embedded), and external artefacts may even be constituent parts of cognitive systems (extended).

Distributed cognition, which has long been incorporated into HCI [14], considers processes in terms of functional relationships between participating elements regardless of where those elements are, and where a cognitive system dynamically employs and coordinates subsystems for a given task. Therefore, cognition involves internal processes, manipulation of objects and the sending/receiving of informational representations among agents or elements. While the approach does not consider external elements to be physical substrates of an individual's actual cognitive system, it does emphasise that we cannot understand cognition or human performance without understanding how people engage with their environment, including technologies, during cognitive tasks. *Embedded cognition* [29] takes a similar stance, proposing that individuals make use of environmental artefacts (*e.g.*, tools) or characteristics (*e.g.*, spatial relationships) as cognitive aids, supports or shortcuts during cognition, but that thought remains inside the head. If a TfM (or TfT) is used as a form of embedded/distributed cognition, it primarily provides support for established memory systems within the brain, making those processes easier or faster, but not necessarily being required for the process to occur at all. The use of the TfM will continue over time as it provides an otherwise unavailable aid to thought: the user cannot internalise the information that is provided. For example a user is planning to revisit a city and plan a route along the best landmarks. The agent in his AR glasses shows a 3D reconstruction of his previous trip/route as a cognitive aid.

The "extended mind" (or *extended cognition*) hypothesis [4], in contrast, takes a more provocative (though potentially less substantiated [13, 29]) perspective. It proposes that, if something in the environment functions or participates in a process in the same way it would if it was done inside the head, then it is indeed part of the cognitive system. Clark and Chalmers [4] provide the example of a man with memory impairment (called Otto) who keeps a memo pad with him at all times in which he keeps important information. They propose that the information in the memo pad acts in the same way as factual information would in the brain of a healthy person - it is deeply integrated into Otto's cognitive processes - and so it is arguably part of his cognitive system. They propose four criteria for determining if a situation counts as "extended cognition": *reliability* (the element is always/reliably present to be used), *accessibility/availability* (information is found easily and quickly), *trust* (the information is trustworthy) and *past endorsement* (the individual has used, approved, or endorsed the information in the past). There are a number of modern equivalents to Otto's pad, including note apps, calendar apps, reminders or AI agents.

Academics have argued that, while the internet and smartphones can provide considerable support for memory - particularly semantic [9, 36], prospective [17] and, to a lesser extent, episodic [5, 30] - these uses should not yet be considered *extended cognition/memory* [35], partly due to a failure to meet the availability criterion (information can take time to find). Smart [35] then proposes that the web could be considered an extended memory (cognitive) system by making online information more accessible and more suitably poised to influence and scaffold our everyday thoughts and actions. It could be argued that genAI and tools for memory could meet these improvements, as AI can retrieve information from a users past and provide it quickly, easily and in ways designed to influence thought (decision-making, understanding, opinion etc).

As genAI improves, particularly LLMs and agentic AI, all four criteria for extended memory may be increasingly met. The person's memories (and associated information) that are presented will be more *reliable* (*i.e.* dependably present through smartphones, AR glasses, AI pins etc), trivially *accessible* (through natural language input and output) and, depending on the accuracy of the models, *trusted* and *endorsed* (*i.e.* user has approved of previous AI outputs). The extent of extended memory may partly depend on how often, in what ways, and in what circumstances a user engages the TfM functionality. For a TfM to constitute extended memory it would need to perform functions that the user may be capable of but that they choose to offload, or those that they may not be currently capable of but that complement some internal process. An example is someone who is discussing the events of a social occasion with their agentic AI who captured (or otherwise understands) what happened (via microphone, cameras with VLMs *etc.*). The user gives their perspective on an awkward conversation while the AI provides an alternative, perhaps more factual, portrayal to help clarify what happened. However, the autonomous nature of AI adds complications. If AI is doing the remembering/thinking on behalf the user without them acting, this becomes less like extended memory and more like potentially maladaptive offloading.

2.2 Scaffolding

Cognitive scaffolding typically refers to a support, structure or aid that is provided to an individual to help them learn something new, often in the context of caregivers and children. The need for the scaffold reduces over time until it is no longer needed. Greenfield [12] described the scaffolding metaphor as having five characteristics: it provides *a support*; it functions as *a tool*; it *extends the range* of the worker; it allows the worker to accomplish a task *not otherwise possible*; and it is *used selectively* to aid the worker where needed. Sutton [40] adds that scaffolding is "merely temporary, to be dispensed with at the appropriate stage of development". Therefore, an external element should be considered a memory or cognitive scaffold if it helps an individual improve memory competence or learning and achieve self-sufficiency, rather than an ongoing tool or one that performs a task the individual is already capable of.

Scaffolding is a useful general concept for understanding the use and influence of a given technology for memory (or cognition). The metaphor is easy to grasp and has simple (though vaguely delimited) criteria for evaluating the given support. However, the term could be inappropriately applied to TfMs (or TfTs) that provide any form of support to a user regardless of how, why and over how long the tool is used, and how/whether it affects the user's memory. In these cases it may gloss over maladaptive use or reliance under the guise of "scaffolding". A number of AI functions could certainly count as scaffolding, wherein they allow a user to remember something they otherwise could not, for example teaching users memory techniques like mnemonics or chunking until they are able to apply them independently. However, other AI functions enact an achievable memory process on behalf of the user, such as telling them where they left their keys, or Meta's Persistent Assistant [3]. Whether 'scaffolding' is the right description will depend on the regularity and timescale of use: temporarily & selectively vs. used repeatedly as a matter of course. Scaffolding is generally considered beneficial, but some TfM may not be scaffolding, or not used like scaffolds, as their use is practical and output driven rather than for the purpose of learning or cognitive improvement. In these cases their impact could be maladaptive or harmful.

2.3 Transactive Memory

Transactive memory describes how groups of people collaboratively store and provide information. One person is not capable of knowing everything, and so there are benefits to social groups where important knowledge or skills are shared between different members. There is also a responsibility on each member to hold and provide useful knowledge. Wegner [44] says that transactive memory systems (TMS) depend on three processes: *Directory Updating* (knowledge of who knows what), *Information Allocation* (processes that determine who is to be responsible for storage) and *Retrieval Coordination* (plans for gaining information from members based on relative expertise). While TMSs have historically been applied to interactions/processes among people, the proliferation of information technologies and internet-connected devices has led some to propose that the internet or smartphones may now act as transactive memory partners [42].

In 2013 Ward [42] referred to the internet as a "*supernormal stimulus*": a disruptive form of transactive memory partner that may lead to maladaptive use, because it is an all-encompassing one-stop shop (which may have replaced our semantic memory [9]). When choosing transactive memory partners (related to *Information Allocation* and *Retrieval Coordination*) we prioritise efficiency, accessibility and expertise, which the internet - and arguably LLMs/agentic AI - excels at [9, 42]. It is important to understand whether TfM/TfT are treated like transactive memory partners by users because they, like the internet, may be designed in a way where they are not a bi-directional partner: they may place no responsibility on the user to remember anything in return. Not only could these disruptive technologies affect motivation and ability to learn, they may also impair *metamemory* (knowing what we know, knowing how to learn), impeding the formation of new memories [42].

Heersmink & Sutton [13] argue that internet-based memory is not a true transactive partner. Unlike in social transactive systems, there is no two-way dialog or negotiation/agreement to meet the *Directory Updating* or *Information Allocation* criteria. They state that "*The main reason for [not considering human-internet interactions as TMS] is that in general these interactions are unidirectional, lacking the more deeply integrated two-way or reciprocal information flow which characterises some socially distributed cognitive systems*" (p.151). However, this could change with genAI TfMs, where they meet more of Wegner's criteria to count as a true transactive partner. LLMs/agents can be conversed with and told or configured what to remember, adding in bi-directional dialog and even negotiation, and it can proactively tell/remind us things based on context or inferred needs. AI tools can have wide-ranging perceived expertise or memory stores and simplify access to a range of complex memories and associated information. AI models also rely on interactions with users for training and improving performance. Wegner *et al.* [43] emphasize the cognitive interdependence of individuals in TMS: all people rely on each other for their cognitive performance. The dialogical nature of LLMs/agents means that memory processes and sharing could move more towards social TMS processes, and therefore much deeper cognitive integration (and reliance).

2.4 Dimensions of Cognitive Integration

While the theories and frameworks described above have their proponents, they have also met with criticism, such as Clark & Chalmers' [4] four criteria for extended cognition being too general, overlapping or not providing robust real-world examples [13, 29, 37]. As mentioned, Heersmink & Sutton [13] dispute whether web-based technologies can be considered true transactive memory partners. While the arguments in favour of memory/cognition being embedded or distributed are compelling, they do not necessarily help researchers understand in what ways environmental artefacts interact with or affect memory or cognition.

Heersmink & Sutton [13] propose that, whether an interaction with technology counts as *e.g.*, "extended mind" matters less than how they interact, and to what extent. Therefore they conceptualize relations between the internet and its users in terms of *cognitive integration*. They propose 7 dimensions of integration:

- Information Flow (one-way, two-way and reciprocal)
- Accessibility (reliability of information access)
- Durability (one-off, repeated, permanent)
- Trust (in veracity of information)
- Procedural Transparency (ease of use, invisibility of tech)
- Informational Transparency (ease of understanding)
- Individualisation (relevance/specificity to user vs generic)
- Transformation (how a cognitive system is transformed)

These dimensions provide researchers with concrete and testable means of measuring the manner and extent of cognitive integration with a given technology. The authors themselves evaluate the dimensions on three websites: Wikipedia, IMDB (movie database) and Google search. They conclude that the web is not deeply cognitively integrated but instead better seen as scaffolding for memory/cognition: Wikipedia and IMDB had low-to-medium integration while Google search had medium-high.

However, based on the dimensions of integration the authors propose, it can be argued that genAI, particularly agentic AI and TfM/TfT, would be *much* more deeply integrated, arguably counting as a truer form of extended cognition. GenAI/TfM/TfT could have reciprocal *Information Flow*, as agents learn from users and users learn from agents. GenAI is increasingly *Accessible*, especially if we are to adopt wearable devices (*e.g.*, AR glasses) driven by AI [11, 15, 24]. Interactions with genAI/TfT would likely have repeated if not permanent *Durability*, as users regularly get productive, informational or entertainment value. *Trust* will depend on the accuracy of the models and utility of the output, but repeated engagement could act to improve models over time. LLMs can potentially provide very high *Procedural Transparency*, as interactions can occur through natural language interactions over voice or text prompts. *Informational Transparency* will depend on a number of factors regarding how information is presented or visualised to the user, but natural language can improve ease of understanding. AI agents and tools already learn from user input history to tailor or refine responses, and the utility of responses will increase as the specificity to the user's situation increases, which calls for high *Individualisation*. Finally, we are yet to fully see whether or how AI may *Transform* cognitive processes, but early research in HCI suggests that cognitive strategies, processes and outcomes are already changing [20, 33, 41, 47].

3 The Role(s) of AI in Memory Processes: What Are We Forgetting?

Theme 1 of the CHI 26 Tools for Thought workshop covers the design and usage of GenAI tools, including the questions *"What are useful framings for the role of GenAI that inspire effective design and usage?"* and *"Can we theoretically ground [usage] strategies...?"*. In this paper we contribute towards answering these questions by discussing several theories that attempt to explain how cognitive processes might be integrated with, and affected by, environmental artefacts including technologies such as TfM or TfT. We focus specifically on memory, however, the theories (excluding transactive memory) apply to cognition as a whole as well as memory.

These theories have been used to explore the design or impact of some technologies on memory formation (and/or cognition) but primarily the internet, and relatively early/basic internet functions at that (*e.g.*, text webpages, Google search, calendars, reminders). They are potentially powerful theories for understanding cognition as a whole but they are yet to be used to understand how other, more modern, technologies might integrate with memory/cognition. As we have argued in each section, GenAI technologies - including LLMs/VLMs and agents - may be uniquely positioned to deeply integrate into memory and wider cognition due to the way(s) in which they can provide information and the potential ease with which users can engage with them. This could lead to significant benefits to memory: the less that a person needs to hold in biological memory the more capacity they may have to dedicate to more complex recollections or information processing. Or the timely and context-relevant 'recall' of information by a tool could improve task performance, understanding or wellbeing. In contrast, poorly (or maliciously) designed TfM/TfT could allow users to offload increasing amounts of valuable memories to AI, or inadvertently settle into an ecosystem of convenience and efficiency, leading to maladaptive use and losses in memories themselves, metamemory and potentially even brain matter.

Therefore, we argue that the HCI community - in particular, attendees of this workshop - need to engage with theories of cognitive integration and properly interrogate the role of GenAI across memory processes. We currently lack the theoretical underpinnings to fully understand the implications of this, nor the empirical approach to measure any short or long-term impact. To define and practice responsible cognitive augmentation, first we must conceptualize cognitive theories of augmentation (what does it change/impact and how), then define how we can empirically examine the human impact of prolonged and everyday usage, to prove there exists irresponsible augmentation. Only then can we identify the dimensions, criteria or characteristics that can be designed for and/or evaluated in the pursuit of *responsible Tools for Memory*.

4 Grand Challenges in Tools for Memory

As in our recent work [45], we recommend that academics, companies and policy makers prioritise the following issues, and we are keen to discuss and collaborate with workshop attendees and interested stakeholders in order to work towards these goals.

Identify Foundational Theories. To better and more quickly understand how memory (and wider cognition) is affected by genAI the community needs to not only identify theories that ground usage strategies but also theories that explain how cognition is integrated with technologies more generally.

Establish the Key Benefits and Risks. Concrete evidence is needed to establish where the main benefits and harms of TfM lie, allowing us to target priority needs or vulnerabilities. This could be achieved through lab studies, participatory design or consultation with domain experts in memory and neuroscience.

Automation and Memory. Memory consolidation benefits from active engagement with events but genAI can perform a number of functions autonomously on behalf of the user, who may prefer 'query-less' interaction. Therefore, it is necessary to understand the effects of automation vs. active engagement, as well as the acceptance or need for usage prompts in the general public.

Long-Term Effects of Use. The consolidation of memory requires time and repetition. If technology has supportive or disruptive effects then they may occur over long time periods and so, alongside lab studies, longitudinal research is also needed.

Acknowledgments

Funded by UK Research and Innovation (UKRI) under the UK's Horizon Europe funding guarantee [EP/Z000068/1] (AUGSOC).

References

[1] Nathaniel Barr, Gordon Pennycook, Jennifer A. Stolz, and Jonathan A. Fugelsang. 2015. The brain in your pocket: Evidence that Smartphones are used to supplant thinking. *Computers in Human Behavior* 48 (2015), 473–480. doi:10.1016/j.chb.2015.02.029

[2] Susan Bluck, Nicole Alea, Tilmann Habermas, and David C. Rubin. 2005. A TALE of Three Functions: The Self–Reported Uses of Autobiographical Memory. *Social Cognition* 23, 1 (2005), 91–117. doi:10.1521/soco.23.1.91.59198 arXiv:https://doi.org/10.1521/soco.23.1.91.59198

[3] Hyunsung Cho, Jacqui Fashimpaur, Naveen Sendhilnathan, Jonathan Browder, David Lindlbauer, Tanya R. Jonker, and Kashyap Todi. 2025. Persistent Assistant: Seamless Everyday AI Interactions via Intent Grounding and Multimodal Feedback. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 59, 19 pages. doi:10.1145/3706598.3714317

[4] Andy Clark and David Chalmers. 1998. The Extended Mind. *Analysis* 58, 1 (1998), 7–19. http://www.jstor.org/stable/3328150

[5] Tilman Dingler, Passant El Agroudy, Rufat Rzayev, Lars Lischke, Tonja Machulla, and Albrecht Schmidt. 2021. *Memory Augmentation Through Lifelogging: Opportunities and Challenges*. Springer International Publishing, Cham, 47–69. doi:10.1007/978-3-030-30457-7_3

[6] Guangheng Dong, Hui Li, and Marc N. Potenza. 2017. Short-Term Internet-Search Training Is Associated with Increased Fractional Anisotropy in the Superior Longitudinal Fasciculus in the Parietal Lobe. *Frontiers in Neuroscience* Volume 11 - 2017 (2017). doi:10.3389/fnins.2017.00372

[7] Guangheng Dong and Marc N. Potenza. 2015. Behavioural and brain responses related to Internet search and memory. *European Journal of Neuroscience* 42, 8 (2015), 2546–2554. doi:10.1111/ejn.13039 arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/ejn.13039

[8] Guangheng Dong and Marc N. Potenza. 2016. Short-term Internet-search practicing modulates brain activity during recollection. *Neuroscience* 335 (2016), 82–90. doi:10.1016/j.neuroscience.2016.08.028

[9] Joseph Firth, John Torous, Brendon Stubbs, Josh A. Firth, Genevieve Z. Steiner, Lee Smith, Mario Alvarez-Jimenez, John Gleeson, Davy Vancampfort, Christopher J. Armitage, and Jerome Sarris. 2019. The "online brain": how the Internet may be changing our cognition. *World Psychiatry* 18, 2 (2019), 119–129. doi:10.1002/wps.20617 arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/wps.20617

[10] Iason Gabriel, Arianna Manzini, Geoff Keeling, Lisa Anne Hendricks, Verena Rieser, Hasan Iqbal, Nenad Tomašev, Ira Ktena, Zachary Kenton, Mikel Rodriguez, Seliem El-Sayed, Sasha Brown, Canfer Akbulut, Andrew Trask, Edward Hughes, A. Stevie Bergman, Renee Shelby, Nahema Marchal, Conor Griffin, Juan Mateos-Garcia, Laura Weidinger, Winnie Street, Benjamin Lange, Alex Ingerman, Alison Lentz, Reed Enger, Andrew Barakat, Victoria Krakovna, John Oliver Siy, Zeb Kurth-Nelson, Amanda McCroskery, Vijay Bolina, Harry Law, Murray Shanahan, Lize Alberts, Borja Balle, Sarah de Haas, Yetunde Ibitoye, Allan Dafoe, Beth Goldberg, Sébastien Krier, Alexander Reese, Sims Witherspoon, Will Hawkins, Maribeth Rauh, Don Wallace, Matija Franklin, Josh A. Goldstein, Joel Lehman, Michael Klenk, Shannon Vallor, Courtney Biles, Meredith Ringel Morris, Helen King, Blaise Agüera y Arcas, William Isaac, and James Manyika. 2024. The Ethics of Advanced AI Assistants. arXiv:2404.16244 [cs.CY] https://arxiv.org/abs/2404.16244

[11] Google. 2025. A new look at how Android XR will bring Gemini to glasses and headsets. https://blog.google/products/android/android-xr-gemini-glasses-headsets/.

[12] Patricia M. Greenfield. 1984. *A theory of the teacher in the learning activities of everyday life.* Harvard University Press.

[13] Richard Heersmink and John Sutton. 2020. Cognition and the web: extended, transactive, or scaffolded? *Erkenntnis* 85, 1 (Feb. 2020), 139–164. doi:10.1007/s10670-018-0022-8

[14] James Hollan, Edwin Hutchins, and David Kirsh. 2000. Distributed cognition: toward a new foundation for human-computer interaction research. *ACM Trans. Comput.-Hum. Interact.* 7, 2 (June 2000), 174–196. doi:10.1145/353485.353487

[15] Axios Ina Fried. 2025. In Meta's AI future, your friends are bots. https://www.axios.com/2025/05/02/meta-zuckerberg-ai-bots-friends-companions.

[16] Matthew Jamieson, Breda Cullen, Marilyn McGee-Lennon, Stephen Brewster, and Jonathan J. Evans. 2014. The efficacy of cognitive prosthetic technology for people with memory impairments: A systematic review and meta-analysis. *Neuropsychological Rehabilitation* 24, 3-4 (2014), 419–444. doi:10.1080/09602011.2013.825632 arXiv:https://doi.org/10.1080/09602011.2013.825632 PMID: 23957379.

[17] Matthew Jamieson, Breda Cullen, Marilyn McGee-Lennon, Stephen Brewster, and Jonathan J. Evans. 2014. The efficacy of cognitive prosthetic technology for people with memory impairments: A systematic review and meta-analysis. *Neuropsychological Rehabilitation* 24, 3-4 (2014), 419–444. doi:10.1080/09602011.2013.825632 arXiv:https://doi.org/10.1080/09602011.2013.825632 PMID: 23957379.

[18] Matthew Jamieson, Brian O'Neill, Breda Cullen, Marilyn Lennon, Stephen Brewster, and Jonathan Evans. 2017. ForgetMeNot: Active Reminder Entry Support for Adults with Acquired Brain Injury. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) *(CHI '17)*. Association for Computing Machinery, New York, NY, USA, 6012–6023. doi:10.1145/3025453.3025888

[19] Veronika Krauß, Mark McGill, Thomas Kosch, Yolanda Maira Thiel, Dominik Schön, and Jan Gugenheimer. 2025. "Create a Fear of Missing Out" - ChatGPT Implements Unsolicited Deceptive Designs in Generated Websites Without Warning. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 857, 20 pages. doi:10.1145/3706598.3713083

[20] Hao-Ping (Hank) Lee, Advait Sarkar, Lev Tankelevitch, Ian Drosos, Sean Rintel, Richard Banks, and Nicholas Wilson. 2025. The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 1121, 22 pages. doi:10.1145/3706598.3713778

[21] Franklin Mingzhe Li, Di Laura Chen, Mingming Fan, and Khai N. Truong. 2019. FMT: A Wearable Camera-Based Object Tracking Memory Aid for Older Adults. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 3, 3, Article 95 (Sept. 2019), 25 pages. doi:10.1145/3351253

[22] Xiaoyue Liu, Xiao Lin, Ming Zheng, Yanbo Hu, Yifan Wang, Lingxiao Wang, Xiaoxia Du, and Guangheng Dong. 2018. Internet Search Alters Intra- and Interregional Synchronization in the Temporal Gyrus. *Frontiers in Psychology* Volume 9 - 2018 (2018). doi:10.3389/fpsyg.2018.00260

[23] Meta. 2025. Introducing Aria Everyday Activities. https://www.projectaria.com/datasets/aea/.

[24] Meta. 2025. Introducing Orion, Our First True Augmented Reality Glasses. https://about.fb.com/news/2024/09/introducing-orion-our-first-true-augmented-reality-glasses/.

[25] Joseph O'Hagan, Graham Wilson, and Mark McGill. 2026. WebPerceptor: An Open Source Chromium Plugin for Real-Time LLM-Based In-Line, In-Browser Re-Writing of Website Content. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '26)*. Association for Computing Machinery, New York, NY, USA.

[26] Pat Pataranutaporn, Chayapatr Archiwaranguprok, Samantha W. T. Chan, Elizabeth Loftus, and Pattie Maes. 2025. Slip Through the Chat: Subtle Injection of False Information in LLM Chatbot Conversations Increases False Memory Formation. In *Proceedings of the 30th International Conference on Intelligent User*

Interfaces (IUI '25). Association for Computing Machinery, New York, NY, USA, 1297–1313. doi:10.1145/3708359.3712112

[27] Pat Pataranutaporn, Chayapatr Archiwaranguprok, Samantha W. T. Chan, Elizabeth Loftus, and Pattie Maes. 2025. Synthetic Human Memories: AI-Edited Images and Videos Can Implant False Memories and Distort Recollection. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 538, 20 pages. doi:10.1145/3706598.3713697

[28] E.F. Risko, M.O. Kelly, P. Patel, and C. Gaspar. 2019. Offloading memory leaves us vulnerable to memory manipulation. *Cognition* 191 (2019), 103954. doi:10.1016/j.cognition.2019.04.023

[29] Robert D. Rupert. 2004. Challenges to the Hypothesis of Extended Cognition. *The Journal of Philosophy* 101, 8 (2004), 389–428. http://www.jstor.org/stable/3655517

[30] Abigail J. Sellen, Andrew Fogg, Mike Aitken, Steve Hodges, Carsten Rother, and Ken Wood. 2007. Do life-logging technologies support memory for the past? an experimental study using sensecam. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) *(CHI '07)*. Association for Computing Machinery, New York, NY, USA, 81–90. doi:10.1145/1240624.1240636

[31] Lawrence Shapiro. 2019. *Embodied Cognition*. Routledge. https://doi.org/10.4324/9781315180380

[32] Sara J. Shettleworth. 2009. *Cognition, Evolution, and Behavior (2nd Ed)*. Oxford University Press.

[33] Prakash Shukla, Phuong Bui, Sean S Levy, Max Kowalski, Ali Baigelenov, and Paul Parsons. 2025. De-skilling, Cognitive Offloading, and Misplaced Responsibilities: Potential Ironies of AI-Assisted Design. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '25)*. Association for Computing Machinery, New York, NY, USA, Article 171, 7 pages. doi:10.1145/3706599.3719931

[34] Anjali Singh, Karan Taneja, Zhitong Guan, and Avijit Ghosh. 2025. Protecting Human Cognition in the Age of AI. arXiv:2502.12447 [cs.CY] https://arxiv.org/abs/2502.12447

[35] Paul R. Smart. 2017. Extended cognition and the Internet: a review of current issues and controversies. *Philosophy & Technology* 30, 3 (September 2017), 357–390. https://eprints.soton.ac.uk/404284/

[36] Betsy Sparrow, Jenny Liu, and Daniel M. Wegner. 2011. Google Effects on Memory: Cognitive Consequences of Having Information at Our Fingertips. *Science* 333, 6043 (2011), 776–778. doi:10.1126/science.1207745 arXiv:https://www.science.org/doi/pdf/10.1126/science.1207745

[37] Kim Sterelny. 2010. Minds: Extended or scaffolded? *Phenomenology and the Cognitive Sciences* 9, 4 (2010), 465–481. doi:10.1007/s11097-010-9174-y

[38] Benjamin C. Storm and Sean M. Stone. 2015. Saving-Enhanced Memory: The Benefits of Saving on the Learning and Remembering of New Information. *Psychological Science* 26, 2 (2015), 182–188. doi:10.1177/0956797614559285 arXiv:https://doi.org/10.1177/0956797614559285 PMID: 25491269.

[39] Benjamin C. Storm, Sean M. Stone, and Aaron S. Benjamin. 2017. Using the Internet to access information inflates future use of the Internet to access other information. *Memory* 25, 6 (2017), 717–723. doi:10.1080/09658211.2016.1210171 arXiv:https://doi.org/10.1080/09658211.2016.1210171 PMID: 27424847.

[40] John Sutton. 2015. Scaffolding Memory: Themes, Taxonomies, Puzzles. In *Contextualizing Human Memory: An Interdisciplinary Approach to Understanding How Individuals and Groups Remember the Past*, Charles Stone and Lucas Bietti (Eds.). Routledge, 187–205.

[41] Lev Tankelevitch, Viktor Kewenig, Auste Simkute, Ava Elizabeth Scott, Advait Sarkar, Abigail Sellen, and Sean Rintel. 2024. The Metacognitive Demands and Opportunities of Generative AI. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 680, 24 pages. doi:10.1145/3613904.3642902

[42] Adrian F. Ward. 2013. Supernormal: How the Internet Is Changing Our Memories and Our Minds. *Psychological Inquiry* 24, 4 (2013), 341–348. doi:10.1080/1047840X.2013.850148 arXiv:https://doi.org/10.1080/1047840X.2013.850148

[43] Daniel M. Wegner. 1987. *Transactive Memory: A Contemporary Analysis of the Group Mind*. Springer New York, New York, NY, 185–208. doi:10.1007/978-1-4612-4634-3_9

[44] Daniel M. Wegner. 1995. A Computer Network Model of Human Transactive Memory. *Social Cognition* 13 (1995), 319–339. https://api.semanticscholar.org/CorpusID:144862841

[45] Graham Wilson, Thomas Goodge, Joseph O'Hagan, and Mark McGill. 2026. Towards Responsible Memory Augmentation: Surveying the Intended Uses of Everyday Augmented Reality and Artificial Intelligence for Memory Processes. In *Proceedings of the Augmented Humans International Conference 2026 (AHs '26)*. Association for Computing Machinery, New York, NY, USA. doi:tbc

[46] Ka-Ho Wong, Christopher D. Anderson, Cecilia Peterson, Erin Bouldin, Lauren Littig, Neeharika Krothapalli, Trieste Francis, Yvonne Kim, Giselle Cucufate, Jonathan Rosand, Kevin Navin Sheth, and Adam de Havenon. 2025. Rising Cognitive Disability as a Public Health Concern Among US Adults. *Neurology* 105, 8 (2025), e214226. doi:10.1212/WNL.0000000000214226 arXiv:https://www.neurology.org/doi/pdf/10.1212/WNL.0000000000214226

[47] Nandi Zhang, Yukang Yan, and Ryo Suzuki. 2025. From Following to Understanding: Investigating the Role of Reflective Prompts in AR-Guided Tasks to Promote User Understanding. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 1227, 18 pages. doi:10.1145/3706598.3713293

[48] Wazeer Deen Zulfikar, Samantha Chan, and Pattie Maes. 2024. Memoro: Using Large Language Models to Realize a Concise Interface for Real-Time Memory Augmentation. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 450, 18 pages. doi:10.1145/3613904.3642450