International Journal of Innovative Technologies in Social Science

e-ISSN: 2544-9435

Operating Publisher
SciFormat Publishing Inc.
ISNI: 0000 0005 1449 8214

2734 17 Avenue SW,
Calgary, Alberta, T3E0A7,
Canada
+15878858911
editorial-office@sciformat.ca

ARTICLE TITLE	MEMORY-SUPPORT TECHNOLOGIES IN THE AGE OF ARTIFICIAL INTELLIGENCE: COGNITIVE OFFLOADING, SOCIETAL IMPLICATIONS, AND CONDITIONS FOR RESPONSIBLE IMPLEMENTATION
DOI	https://doi.org/10.31435/ijitss.2(50).2026.5257
RECEIVED	22 February 2026
ACCEPTED	16 April 2026
PUBLISHED	30 April 2026
LICENSE	The article is licensed under a **Creative Commons Attribution 4.0 International License.**

MEMORY-SUPPORT TECHNOLOGIES IN THE AGE OF ARTIFICIAL INTELLIGENCE: COGNITIVE OFFLOADING, SOCIETAL IMPLICATIONS, AND CONDITIONS FOR RESPONSIBLE IMPLEMENTATION

Jakub Winiarczyk *(Corresponding Author, Email: jakub.winiarczyk@onet.eu)*
Independent Public Health Care Complex in Pruszków, Pruszków, Poland
ORCID ID: 0009-0008-8163-5580

Karolina Dąbrowska
Wojewódzki Szpital w Łomży, Łomża, Poland
ORCID ID: 0009-0002-9053-2002

Emilia Lenkiewicz
National Medical Institute of the Ministry of the Interior and Administration, Warsaw, Poland
ORCID ID: 0009-0001-4822-8157

Julia Żak
Czerniakowski Hospital, Warsaw, Poland
ORCID ID: 0009-0008-3383-8025

Zuzanna Rybka
University Clinical Center of the Medical University of Warsaw, Warsaw, Poland
ORCID ID: 0009-0001-4056-8800

Weronika Trynkiewicz
Ludwik Rydygier Specialist Hospital, Kraków, Poland
ORCID ID: 0009-0008-2267-7166

Maja Kaczor
National Medical Institute of the Ministry of the Interior and Administration, Warsaw, Poland
ORCID ID: 0009-0001-8218-5429

Alicja Maciejewska
Independent Public Health Care Complex in Pruszków, Pruszków, Poland
ORCID ID: 0009-0009-6903-2434

Marta Omiecińska
Międzyleski Specialist Hospital of Warsaw, Warsaw, Poland
ORCID ID: 0009-0002-3134-8141

Monika Stępińska
University Clinical Center of the Medical University of Warsaw, Warsaw, Poland
ORCID ID: 0009-0008-0347-2704

ABSTRACT

Research objective: This integrative review examines the impact of memory-aided technology on the encoding, storage, retrieval, and coordination of everyday activities and shared information practices: smartphone reminders, cloud storage, wearable lifelogs, clinical cognitive prostheses, and AI personal assistants, as well as the implications that exist at both the individual and institutional levels.
Methodology: An integrative review was conducted by synthesizing evidence across disciplines including Cognitive Psychology, Human-Computer Interaction, Neuro-Rehabilitation, Governance of Artificial Intelligence, and Digital Rights. The data were then organized using a Taxonomy of Evidence-Informed Technology Families.
Findings: Externalizing memory has been found to have numerous benefits, including decreasing cognitive load, improving performance, and increasing independence. However, there are also trade-offs associated with externalizing memory, such as the shift towards "where-to-find" encoding, the potential for altered metacognitive calibration, decreased effort, attentional costs, or dependence on the device.
Societal Findings: These types of devices have also enabled coordination, continuity, and accountability; however, they also pose risks to users' privacy and consent, the use of biased algorithms, users' autonomy, and equal access to these devices.
Conclusion: To responsibly deploy memory-aided technology will require designers to create user-centered and clinically grounded products, evaluate them over time to determine their effects beyond accuracy, incorporate privacy-by-design principles into the development process, create transparent and auditable pipelines for artificial intelligence systems, plan for future compatibility and continuity, provide users with training and support services, and develop regulatory frameworks and standards that protect users' rights.

KEYWORDS

Memory-Support Technologies, Cognitive Offloading, AI-Enabled Memory Assistants, Lifelogging, Clinical Cognitive Prosthetics, Algorithmic Accountability

CITATION

Jakub Winiarczyk, Karolina Dąbrowska, Emilia Lenkiewicz, Julia Żak, Zuzanna Rybka, Weronika Trynkiewicz, Maja Kaczor, Alicja Maciejewska, Marta Omiecińska, Monika Stępińska. (2026) Memory-Support Technologies in the Age of Artificial Intelligence: Cognitive Offloading, Societal Implications, and Conditions for Responsible Implementation. *International Journal of Innovative Technologies in Social Science*. 2(50). doi: 10.31435/ijitss.2(50).2026.5257

1. Introduction

1.1 Human Memory as Socio-Technical Achievement

Human memory is commonly understood as an internal biological system that encodes, stores, and retrieves information. However, most people's experience of remembering involves significant reliance upon external aids and social scaffolds (e.g., notes, calendars, photographs, routines, documentation) (Clark & Chalmers, 1998). Digital technologies have further augmented the external layer of memory-dependent achievement by providing a means for rapidly capturing information, storing information in searchable formats, and retrieving stored information almost instantly across multiple devices and platforms (Firth et al., 2019; Sparrow et al., 2011). Many memory-dependent achievements – meeting obligations, coordinating care, and transferring responsibilities in workplace settings – rely upon hybrid systems that integrate internal cognition with external artifacts and social scaffolding (Clark & Chalmers, 1998). Cognitive scientists refer to such behaviors as cognitive offloading: the process of utilizing physical or digital tools to decrease the internal cognitive load associated with performing tasks and distribute cognitive work to the external environment (Risko & Gilbert, 2016). Cognitive offloading can be both productive and beneficial, especially for prospective memory (the ability to remember to perform a task at a future date) and managing complex goals (Risko & Gilbert, 2016; Scarampi & Gilbert, 2020). Cognitive offloading is, however, social and cultural in nature beyond the productivity aspects, as it has the potential to influence the types of things that individuals encode, how they focus their attention, and what is perceived as competent in educational and occupational settings (Firth et al., 2019; Sachdeva & Gilbert, 2020; Ward et al., 2017).

1.2 AI and Memory Support Technologies: An Inflection Point Between Storage and Proactive Assistance

AI transforms memory support technologies in two ways. First, AI converts passive repositories of information into semantic systems capable of being queried by users via natural language, rather than requiring the users to navigate file hierarchies (National Institute of Standards and Technology [NIST], 2024). Second, AI provides proactive support for users, allowing systems to draw contextual inferences and automatically remind users, summarize events or decisions, and make decisions for users based on contextual inferences, even if users do not request assistance (NIST, 2024). As a result, AI-based memory support technologies create the potential to significantly reduce friction and promote continuity in high demand contexts (e.g., healthcare delivery, workplaces, and education) (NIST, 2024). However, AI-based memory support technologies also create new avenues for error (e.g., opaque selection of retrieved information, biased decision making against certain demographic groups, languages, etc.) (Buolamwini & Gebru, 2018; NIST, 2024), and AI-generated summaries and recommendations that confidently assert incorrect information that may rewrite a person's history (NIST, 2024). Finally, AI-based memory support technologies create accountability challenges related to errors or adverse outcomes, as there may be limited transparency into how AI-based systems reached their conclusions and made recommendations (de Fine Licht & de Fine Licht, 2020; Diakopoulos, 2015; Raji et al., 2020).

1.3 Why Memory Support Technology is a Technology-Society Issue

Memory support technologies are increasingly functioning as infrastructural systems. Here, “infrastructural” refers to systems that become embedded in routine organizational and interpersonal workflows, enabling continuity, coordination, and accountability across actors. That is, memory support technologies are increasingly becoming essential systems that coordinate (via task systems, reminders, etc.), continue (via handover documents, shared calendars, etc.) and hold accountable (via digital records, etc.) various stakeholders and their actions. At the same time, memory support technologies are raising traditional technology-society issues: Who owns and controls personal data collected and stored by memory support technologies? Who has access to that data? How easily can users delete or correct that data? And how will memory support technologies be used to allocate opportunities, and/or impose discipline? (Nissenbaum, 2009). When memory support technologies engage in lifelogging or ambient capture, bystander privacy and consent issues come to the fore (Hodges et al., 2006; Sellen et al., 2007; Sellen & Whittaker, 2010). When memory support technologies leverage AI, new governance issues emerge: auditability, documentation, contestability, compliance with regulations and standards, and the like (European Parliament and Council of the European Union, 2024; Felzmann et al., 2019; NIST, 2023; Organisation for Economic Co-operation and Development [OECD], 2019; UNESCO, 2021).

For this reason, this review is located at the juncture of technology and society. Therefore, it will synthesize cognitive and clinical evidence of the advantages and disadvantages of memory support technologies, as well as translate governance research into actionable system design requirements.

1.4 Research Questions

RQ1: Do memory support technologies impact individual cognition, behavior, and well-being in daily life, clinical settings, and aging-related contexts?

RQ2: Do memory support technologies lead to changes in society and institutions, including new risks, power imbalances, and inequities?

RQ3: What system design conditions (governance, infrastructure, standards, professional practices) are needed to ensure that memory support technologies are implemented in an ethical, effective, and sustainable manner in the era of artificial intelligence?

1.5 Conceptual Lenses for Synthesis

To facilitate interpretation of evidence from disparate disciplines, we employ three conceptual lenses to analyze our findings:

I. Extended and Distributed Cognition: Emphasizes that tools can become functionally incorporated into cognitive systems; thus, memory support is not simply an external aid to cognition, but is an integral component of the cognitive ecology (Clark & Chalmers, 1998).

II. Contextual Integrity: Stresses that privacy and legitimacy depend upon contextually appropriate information flow, and not the abstract distinction between public and private information (Nissenbaum, 2009).

III. Algorithmic Accountability and Auditing: Recognizes that harm typically arises from the lack of transparency in socio-technical processes and that accountability requires documentation of the entire lifecycle of the system, monitoring, and governance roles (Diakopoulos, 2015; Raji et al., 2020).

1.6 Contribution

The purpose of this literature review is to compare evidence-based tradeoffs among families of memory support technologies; identify a set of minimum design requirements to ensure that memory support technologies are implemented in a responsible manner; and outline an applied research agenda to guide incrementally innovative socially-responsible AI-mediated memory support technologies.

2. Methodology

2.1 Design and Rationale

We conducted an integrative review to provide a synthesis of the heterogeneous body of evidence spanning the domains of cognitive psychology, Human Computer Interaction (HCI), Neurorehabilitation, and AI Governance. Integrative Reviews are best suited for a field that encompasses various methodologies and where the goal is to bridge mechanism, outcome, and implementation condition rather than compute a single pooled effect size.

2.2 Data Sources and Search Strategy

We found relevant studies using a structured search methodology in multidisciplinary and discipline-specific databases (i.e., Scopus, Web of Science, PubMed, PsycINFO, IEEE Xplore, ACM Digital Library) and supplemented this with both backward and forward citation tracking. Our keywords to search for were a combination of the idea of memory support through cognitive offloading using AI/Governance as follows: (“cognitive offloading” or “external memory” or “memory aid*” or “prospective memory” or “lifelog*” or “SenseCam” or “wearable camera” or “memory support system”) & (“AI” or “machine learning” or “algorithmic” or “audit*” or “transparency” or “governance” or “ethics” or “rights” or “privacy”).

Our approach to finding research on cognitive offloading was based on a broad definition of cognitive offloading as an interdisciplinary concept whose evaluation can be different depending on how it is defined in each discipline; therefore, we searched with broad search terms and limited our inclusion to the full-text articles. As such, our literature searches were for literature that has been published between 1998 and 2024. Finally, all literature included in our analysis was written in English.

2.3 Eligibility Criteria

Criteria for Inclusion:

I. Empirical Studies and Systematic Reviews Evaluating Memory-Support Technologies for Both Healthy Users and Clinical and Aging Populations;

II. HCI Research on Digital/Wearable Memory Systems or Personal Information Management That Have Explicit Memory-Support Outcomes;

III. Literature on Governance/Rights and High-Authority Standards/Policies Directly Influencing Real-World Deployment of AI-Mediated Memory Support (i.e., Accountability Frameworks, Audit Practices, Risk Management Guidance, Binding Regulation).

Criteria for Exclusion:

I. Non-Scholarly Commentary Without Empirical or Normative Grounding;

II. Cognitive Training Research Without an External Memory-Support Component;

III. Purely Technical AI Papers Without Relevance to Memory Support.

2.4 Screening and Reporting Transparency

All records obtained from the database searches as outlined in Section 2.2 were imported into a reference management file, and duplicates were removed prior to screening. Study selection involved a two-step process. First, title and abstract screening was conducted against the eligibility criteria (Section 2.3) to eliminate obviously irrelevant records for example, purely technical AI papers that lack a memory-support element, cognitive training studies that do not use an external aid, or non-academic commentaries without empirical or normative support. Second, full texts of possibly relevant records were evaluated based on the eligibility criteria and retained only when they provided direct evidence or high-authority normative guidance related to memory-support technologies across the review's three analytical levels: individual cognitive outcomes; organizational/interpersonal coordination; societal/system governance outcomes.

Database retrieval was supplemented by backward and forward citation tracking to find influential foundational work and cross-disciplinary evidence that is not consistently indexed under a single keyword set in interdisciplinary databases. As the review is situated at the intersection of technology and society, the screening process also included high-authority standards and regulatory and policy documents which directly influence how AI-mediated memory support is implemented in practice (NIST risk-management guidelines for generative AI, OECD AI Recommendation, UNESCO AI ethics recommendations, and the EU AI Act) (European Parliament and Council of the European Union, 2024; NIST, 2024; OECD, 2019; UNESCO, 2021) and therefore meet the inclusion criteria for governance and digital rights sources (Section 2.3).

The final dataset contained 34 sources ranging in publication date from 1998–2024 consisting of peer-reviewed empirical research, systematic reviews, HCI conference proceedings, and high-authority governance instruments. These included sources constitute the reference list of this review. To enhance reporting transparency, all included sources were retained in a reference management file, and inclusion decisions were documented at the full-text stage. Given the interdisciplinary scope and the inclusion of high-authority governance instruments alongside empirical studies, we report the final included set (n = 34) rather than a formal PRISMA flow diagram.

2.5 Data Extraction and Synthesis

For each of the included items, we extracted: Technology Type; Memory Function Targeted (Prospective, Episodic/Autobiographical, Semantic, Everyday Functional Memory); Population and Setting; Study Design; Primary Outcomes and Measures; Reported Benefits and Harms; Implementation Requirements (Training, Support Services, Infrastructure, Data Governance, Oversight). Then we completed a Thematic Synthesis and Comparative Mapping Across Three Levels:

Level 1: Individual Outcomes (Cognitive Performance, Attention, Metacognition, Well-being).

Level 2: Organizational/Interpersonal Outcomes (Coordination, Handovers, Caregiving Routines, Documentation).

Level 3: Societal/System Outcomes (Rights, Governance, Inequality, Infrastructure, Standards).

2.6 Evidence Strength and Limitations

The strength of the evidence varies across technology families. The majority of cognitive offloading research is based on controlled laboratory experiments with healthy participants; the majority of clinical assistive technology is comprised of small trials, single-case designs, and practical evaluations; the majority of governance sources are peer-reviewed conceptual works, standards, and regulation. Therefore, we have utilized a tripartite methodology to provide evidence and caution against making causal claims particularly regarding long-term cognitive effects.

2.7 Ethics and Governance Sensitivity Screening

Because many memory-support technologies are designed to collect intimate personal data, we have treated ethics and governance as first-order extraction categories as opposed to afterthoughts. For each source, we documented if the technology involved passive collection of data, if bystanders could be recorded, which data types were collected (text, images, location, health-related cues), and if deletion, access control, or portability was mentioned. We also documented if studies reported adverse events or user distress, and if support services were included (training, caregiver involvement, clinical monitoring).

During the synthesis process, we compared the extraction fields across technology families to examine differences in implementation burden. This does not constitute formal risk assessment, but links cognitive outcomes to the institutional conditions that will determine whether a technology can be responsibly deployed for IJITSS readers.

3. Results

3.1 Comparative taxonomy

We find four technology families that aggregate at the evidence base:

T1 Everyday digital offloading (search, saving, reminders). T2 Lifelogging capture-and-replay (wearable cameras, passive recording, diaries). T3 Clinical cognitive prosthetics (structured notebook/calendars; electronic aids that are embedded in rehabilitation and care routines). T4 AI-enabled memory assistants (context-aware retrieval, semantic search over personal data, summarization, proactive prompting, cross-modal retrieval).

They differ in terms of the memory functions they target, their implementation burdens and governance risks.

Table 1. Evidence-informed taxonomy of memory-support technologies and comparative trade-offs

Technology family	Core functions & targeted memory	Benefits supported by evidence	Risks & implementation notes (incl. governance)
T1 Everyday digital offloading (search, saving, reminders)	Information lookup; task/deadline management; learning. Mainly prospective & semantic memory.	Improved task completion and strategic resource allocation; reduced interference when saving information externally.	May disfavour encoding towards "where-to-find" knowledge; can encourage effort-minimization and introduce attentional costs; platform dependence and exposure of privacy.
T2 Lifelogging capture-and-replay (wearables, diaries)	Passive capture and then later replay. Episodic/autobiographical memory via retrieval cues.	Richer retrieval cues; psychosocial benefits when embedded in care in some dementia/MCI contexts.	High privacy and consent burden (including on bystanders), security risks, and potential emotional harms in replay.
T3 Clinical cognitive prosthetics (structured notebooks, electronic aids)	Providing rehabilitation after TBI/ABI; supporting independence; adherence. Mainly prospective and functional everyday memory.	Functional improvements when training/support are given; evidence for structured, coached adoption.	Usability barriers and abandonment in the absence of support; requires service model (selection, training, monitoring) and integration into care routines.
T4 AI-enabled memory assistants (context-aware selective retrieval, summarization)	Proactive prompting & semantic search over personal data. Summarization, decision support, and cross-modal retrieval.	Potential to personalize support and increase adoption; supports organizational continuity & handovers.	Opacity, bias, and surveillance; accountability gaps; rights conflicts (deletion, contestability, explanation). Require auditable pipelines and authority/robust governance.

3.2 T1: Everyday digital offloading—cognitive enhancement and trade-offs

3.2.1 Performance gains through strategic redistribution

Common to studies on digital offloading is the finding that external supports reliably enhance performance by lowering internal demands on memory, especially for prospective memory tasks and complex goal management. Successful offloading can free up internal cognitive resources that can be reallocated to other operations, leading to successful task completion with reduced errors when heavily loaded (Risko & Gilbert, 2016). The successful storage of potentially interference-prone information externally can also protect against interference and can lead to greater overall learning of new information consistent with strategic theories of offloading where external storage acts as a control mechanism over limited internal resources (Storm & Stone, 2015).

3.2.2 Encoding shifts: remembering where instead of what

Of course, successful offloading does not merely provide extra bandwidth, but instead also alters what is encoded in the mind at all. Research surrounding the "Google effect" suggests that when people expect to be able to find information later, they tend to remember where to find it, and the route through information, rather than the information itself (Sparrow et al., 2011). This may be adaptive (in a networked environment), but has implications for education and on-the-job training: we may increasingly judge competence by how well people can find what they need, rather than internalize it.

3.2.3 Attentional costs and presence of device

Digital memory support can also have background costs. Experimental evidence suggests that merely the presence of one's phone can reduce available cognitive capacity when performing difficult tasks, even when not using the phone (Ward et al., 2017). This suggests a design imperative: memory supports should intrude less on attention (e.g., less interruption, batching, user control over timing of notifications).

3.2.4 Metacognitive calibration and overreliance

The decisions people make about memory offloading are shaped by the metacognitive and effort trade-off. Experimental evidence suggests people sometimes set more reminders than they optimally should, perhaps due to underconfidence and achieving minimum effort; this can lead to use becoming self-reinforcing (Sachdeva & Gilbert, 2020; Scarampi & Gilbert, 2020). Thus, for AI assistants, simply decreasing friction of offloading can lead to dependence, unless they also help support calibration (e.g., explaining why a reminder was suggested, thresholds for remembering without reminders, etc., when stakes or learning goals suggest that this is a good idea).

3.2.5 Everyday practices of "digital prosthetic memory" in the wild

Laboratory studies can show us mechanism-level effects, but HCI reminds us that even everyday offloading is shaped by social routines, tool ecosystems, and breakdown recovery. In a classic CHI study, Kalnikaitė and Whittaker (2007) used mixed methods approaches to look at when and why people used digital systems as "prosthetic memory," finding that users offloaded selectively according to perceptions of future value, searchability, and social cost of forgetting. The act itself of choosing to offload information was tied up in negotiations with others—externalizing commitments to demonstrate reliability, to provide accountability in group work, and to create shared references in family and team life. This supports a tech–society ridealong interpretation: external memory is more than a cognitive shortcut, it's also social signaling and coordination. The same study also underlines a design implication relevant to AI assistants: reliability and recoverability in use matter to users. Predictability through searchability, resilience to failure in use leads to trust from users; they may disengage when retrieval fails, or it becomes more effortful to offload than to remember.

Designers should treat retrieval as an interaction that is socio-technical: look at interfaces for "remembering by reconstruction," allow search cues on time/place/people/topic, make failure states legible so that users can see these pathways and "fix" the failure without loss of confidence.

3.2.6 Photo-taking, encoding depth, and the substitution effect

Not all forms of offloading promote later remembering to the same degree. Evidence on photo-taking suggests a potential substitution effect where the individual starts to treat the photo-taking device as the repository, but may encode less depth internally as a result. In a museum study, Henkel (2014) showed how taking photos resulted in poorer memory for actually viewed objects vs. simply observing them—suggesting that capture can actually displace attending. This implication matters for lifelogging and AI summarization, in that if capture or an AI's summary reduces active engagement and encoding depth, then longer term learning and memory richness in an autobiography could suffer even if the episodic retrieval is stronger. So responsible implementation should gauge not only "can users retrieve later?" but also "does this system change how they attend, encode and learn today?"

3.3 T2: Lifelogging capture-and-replay—rich cues, psychosocial effects, and privacy burdens

3.3.1 Mechanism: retrieval cues and autobiographical continuity

T2 systems that "capture-and-replay" dense traces of everyday life (e.g., images and other sensor data) from which the system later serves as a cue for episodic/autobiographical memory retrieval. SenseCam research indicates that these passively captured images may cue memory otherwise inaccessible to conscious retrieval, offering an alternative route towards reconstructing the autobiographical narrative (Hodges et al., 2006; Sellen et al., 2007). Within a systematic review synthesis, wearable cameras proved useful tools for investigating and remediating autobiographical memory impairments although the outcomes differed across populations and implementation styles (Allé et al., 2017; Silva et al., 2018).

3.3.2 Clinical and aging contexts: beyond memory to well-being

In dementia-related contexts the capture-and-replay approach is investigated for use in rehabilitation therapies for its potentially positive influence on well-being and sense of subjective cognitive functioning in addition to the more straightforward performance on memory tests (Silva et al., 2017). While this appears to provide evidence of therapeutic value even for severely impaired memory systems, caveats point out the need for close supervision due to the idiosyncratic case-based impact of this system (Pauly-Takacs et al., 2011).

3.3.3 Governance risks: consent, bystanders, security and emotional harms

Privacy and consent burdens are intensified by the total capture nature of the lifelogging practice, a movement towards capturing everything going on around us. The critique of the total capture ideal argues that it is psychologically and ethically problematic, thus determining an acceptable form of total capture becomes a matter of making these systems selective and user-controllable (Sellen & Whittaker, 2010). Governance risks include access to a life-trace archive stored without authorization; function creep, where the system is used for secondary uses other than the intention for using it; and emotional harms where replaying materials may trigger distress or rumination. These risks are amplified in certain populations and care settings involving power imbalances that may diminish meaningful voluntariness.

3.4 T3: Clinical cognitive prosthetics—effectiveness of the modality correlates with service models

3.4.1 Structured, coached adoption linked to effectiveness

Clinical cognitive prosthetics are typically a structured notebook or calendar, electronic prompts or reminders, or rehabilitation task procedures following memory guides.

Systematic evidence supports that these technologies can assist people with memory impairments in everyday activities, but training and support from a helper or coach are necessary to ensure this is effective (Jamieson et al., 2014). This emphasises the relationship between usability and learning: absentee onboarding tends to lead to drop-outs as users fail to engage with the platform (Boman et al., 2007).

3.4.2 Mild cognitive impairment: Memory Support System (MSS)

A randomized trial of a structured calendar/notebook showed functional benefits for people with mild cognitive impairment when participants and their partners are trained on the use of this Memory Support System (MSS) (Greenaway et al., 2013). This also reveals a social mechanism: effective memory support is informal/life routine-based, and so mediated through caregivers and partners employing and practicing the rules-of-thumb and remembering plans.

3.4.3 Traumatic brain injury, Acquired brain injury: matching to routine

Systematic reviews suggest that electronic assistive technology following traumatic brain injury as well as acquired brain injury (ABI); effectiveness is contingent on appropriate matching of tools typically to phase of targeting (or overlapping encoding, storage, or retrieval), as well as the ongoing availability of technical and human support. Users also tend toward combining both technological and non-technological strategies and that uptake correlates with expected benefit, usability of the technology, and support present within the environment (Jamieson et al., 2017). In practice, successful deployment looks like a service: the right tools for people's needs fit into their existing workflows via selection, training, monitoring, and iteration, and then hopefully be seamlessly integrated into existing care workflows.

3.5 T4: AI-enabled memory assistants—personalization plus new failure modes

3.5.1 Potential benefits for continuity and coordination

By extending this logic, AI-enabled memory assistants build on offloading by enabling semantic search, and then proactive functions over personal archives: less friction finding something, continuity and coordination across work/care hand-off, and proactive prompting of prospective memory. Where cognitive demands are already high, cost of an error is high (healthcare adherence, caregiving, complex work coordination), this seems like it could be particularly valuable (NIST, 2024).

3.5.2 Distinct risks: misremembering by synthesis, opacity, and bias

The ability of AI assistants to generate summaries, infer what is important based on user-provided data, and combine multiple data streams can introduce distinct failure modes: plausible yet incorrect summaries ("misremembering by synthesis"), context collapse (surfacing private information at inappropriate times), and differential performance across user groups defined by language, accent, disability, and demographic characteristics (NIST, 2023, 2024; Nissenbaum, 2009; Mitchell et al., 2019; Raji et al., 2020). Additionally, because many of these systems operate as black-box systems, it may be difficult to determine which data sources contributed to a system output and why a particular item was resurfaced (Felzmann et al., 2019; de Fine Licht & de Fine Licht, 2020). Evidence of demographic disparities in commercial facial analysis accuracy demonstrates how multimodal AI can perform unevenly across subgroups, motivating bias evaluation for systems that index personal images or infer identity attributes (Buolamwini & Gebru, 2018; Gebru et al., 2018).

3.5.3 Accountability and auditability as implementation prerequisites

Accountability is an important part of governance when it comes to algorithms because there are potential harms associated with both the process of how decisions are made, as well as the output of those decisions. Frameworks related to algorithmic accountability suggest that decision-making processes should be traceable; therefore, a level of responsibility needs to be assigned (Diakopoulos, 2015). Auditing frameworks for internal algorithmic use propose documenting and monitoring all aspects of an AI system's lifecycle from data to model development, in addition to assigning roles throughout (Raji et al., 2020). When considering the support of memory, this suggests that robust provenance tracking should capture what data were used, what transformations occurred, what confidence was associated with outputs, and how users can correct errors (Raji et al., 2020).

3.5.4 Documentation practices for accountable deployment

There are ways of putting accountability mechanisms into practice that are bolstered by documentation practices, datasheets for datasets and model cards, that make explicit intended use cases, and limitations, subgroup performance, and known risks (Gebru et al., 2018; Mitchell et al., 2019). Documentation for our AI memory assistants might look like: what data sources are ingested; what are the rules for retention and deletion; what are the limits to rate and fidelity of summarization; what are protection protocols to stop other people from being able to infer sensitive attributes from my memory; and what other procedures exist for incident response?

3.6 Cross-family synthesis: common mechanisms and trade-offs

Across T1–T4, there is a recurring pattern: functional benefits are seen, but they are dependent on the quality of implementation and the surrounding socio-technical context.

Five cross-cutting trade-offs were identified as follows:

V. Offloading/Dependence: The greater the ability to offload, the greater the potential for dependence on others to provide that support. This could result in less calibrated metacognition (the ability to evaluate how well you remember things) and confidence in recalling them without assistance.

II. Richness of Personal Traces (Retrieval Cues)/Privacy/Consent Burden: Rich personal traces could enhance retrieval, but the richer traces become the greater privacy concerns and the more complex it becomes to obtain meaningful consent for bystanders.

III. Convenience/Attention Fragmentation/Superficial Processing: Supporting retrieval by always being able to capture or prompt the user can facilitate memory recall, but it can cause attention fragmentation and shallower processing which has consequences for learning and memory richness.

IV. Personalization/Opaque Processes/Bias: AI based systems can personalize and selectively resurface the past (summarize) to better support retrieval; however, the processes involved in personalization and summarization can be opaque and therefore subject to varying performance levels when applied to different user populations, with subsequent impact on detection/correction of errors.

V. Continuity/Accountability/Risk of Surveillance: Continuous records can provide continuity of care/work and/or support users through documentation; however, continuous records can also serve as a mechanism for surveillance and secondary use (depending on the governing body and mechanisms to control access).

3.7 System-level conditions for ethical and sustainable implementation

The evidence suggests that the successful and safe implementation of memory-support technologies relies on system-level conditions related to collective infrastructure, rather than exclusively user choices. We define six interlocking conditions (Table 2) that collectively constitute the minimum viable requirements for responsible implementation.

Table 2. System-level implementation conditions and practical requirements

Condition	What it requires	Why it matters
Governance & accountability	Defined roles (owner, auditor, clinician/case manager where relevant); documentation; monitoring; incident response	Prevents gaps in accountability, and safer iteration and redress.
Data protection & privacy-by-design	Data minimization; meaningful consent; secure storage; deletion/portability; access controls	Minimizes risk of surveillance and misuse; enables redress of rights violations and trust recovery.
Interoperability & infrastructure	Standards- and API-based formats, connections into care/work systems, reliable connectivity and backups	Enables continuity of care, avoids lock-in to a vendor, and equitable access to services.
Evaluation & assurance	Pre-deployment evaluation for negative side effect spotting; usability testing; ensured longitudinal follow up; reporting on outcome data	Ensures harms are caught earlier; validates real-world benefit beyond lab performance.
Capacity-building	Training for users and professionals; support services; digital literacy programs	Helps adoption and avoid abandonment; reduces dropout.
Ethical and legal alignment	Risk classification; safeguards for vulnerable and at-risk people; transparency and contestability mechanisms	Confirms legitimacy and compliance; protects autonomy and mental privacy.

3.8 Evaluation and assurance: what to measure and what not to ignore

A common risk in the field is an obsession with short-term performance over life impact. To help promote responsible deployment, evaluation efforts must be made to plan for functional and social outcomes, not just the accuracy of retrieval. Table 3 shows a practical evaluation-and-safeguards matrix that can also be rendered as a pre-deployment checklist and as a report based on submitted outcome data.

Table 3. Evaluation metrics and safeguards by technology family

Technology family	Core outcome metrics	Safety / governance metrics	Minimum safeguards
T1 Offloading tools	Task completion; prospective memory failures; cognitive load; learning retention	Attention disruption; overuse of reminders; data sharing footprint	User control of notifications; transparency of storage; export/backup options.
T2 Lifelogging	Autobiographical recall; well-being; caregiver burden	Consent quality; bystander capture frequency; security incidents; emotional distress risk	Selective capture; access controls; redaction; guided replay protocols.
T3 Clinical prosthetics	Functional independence; adherence; quality of life; sustained use	Usability barriers; abandonment; inequity in access to training	Service model: onboarding, training, monitoring; caregiver involvement; accessibility design.
T4 AI assistants	Retrieval relevance; user satisfaction; continuity (handover quality, e.g. from AI to human); more time saved	Hallucination/misremembering rate; subgroup performance; explainability used instead of displacing use-cases; review/audit findings.	Provenance tracking; contestability/correction; model & data documentation; remediations for incidents.

4. Discussion

4.1 Memory support as cognitive ecology and social infrastructure

Our results support a view in which “memory support” becomes a cognitive ecology—the distribution of remembering across tools, routines, and social relations. This helps us justify why “technology alone” is insufficient in clinical contexts, and why governance becomes the focus when memory support scales. In workplaces and education, digital memory infrastructures can improve coordination, but also set new expectations (increasing expectations of availability and documentation) and new forms of discipline (thus enabling metric-based monitoring and potentially surveillance-like governance). In care contexts, memory support can enhance independence, but can shift responsibilities to caregivers and create risks to data for vulnerable individuals.

4.2 Autonomy, calibration, and the danger of dependency

A core social-science question is what kind of autonomy such memory-support tools engender. If one only functions with constant scaffolding, autonomy may be instrumentally improved while self-efficacy declines. Offloading research suggests that reminder-use can be excessive and self-reinforcing (Sachdeva & Gilbert, 2020; Scarampi & Gilbert, 2020). AI-enabled assistants may therefore exacerbate this dilemma, because they both remove friction and can be proactive. Responsible design should therefore help support “calibration”: communicate uncertainty, explain prompting logic, allow users to set thresholds, and avoid making explicit external support the default for all tasks. In pedagogical contexts, calibration can also help protect learning: offloading may be key for logistics but not for concepts requiring internal consolidation.

4.3 From explainability to contestability and redress

While explanations are often a proposed solution to the challenge of opacity, in actuality users often need contestability: the ability to challenge and amend an AI-mediated memory trace created for them. As accountability research highlights, if processes are untraceable and roles ambiguous, they create an opening for responsibility gaps to emerge (Diakopoulos, 2015; Raji et al., 2020). For memory assistants, contestability requires I. provenance: what sources were used to generate a summary? II. correction mechanisms: how are user edits made and persisted across interactions and devices? and III. governance: who owns the memory systems and how are users protected if an incomplete summary causes harm? Transparency is also contextual: what counts as a sufficient explanation is contingent on user role, the stakes at play, and existing legal norms; "more disclosure" can remain unusable or misguided if it is not interpretable (Felzmann et al., 2019). In institutional settings, explanations provided by AI can shape received legitimacy and trust, particularly when systems may impact individuals' access to services and the restrictions they face (de Fine Licht & de Fine Licht, 2020).

4.4 Privacy-by-design and bystander rights as non-negotiable constraints

Lifelogging and ambient capture are not ethically neutral "memory aids". Rather, they are forms of recording; a backup or diary of someone else's past that gains power and constituency by virtue of how the subject's past is recorded. Whether data are 'legitimate' depends on whether their social flow is context-appropriate, not abstractly, nor can it rely on user consent (Nissenbaum, 2009). In practice, privacy-by-design entails user agency through minimization, local processing where appropriate, encryption, fine-grained access, along with deletion/portability commitments. For lifelogging specifically, providing users with selective capture and redaction capabilities can help mitigate unwanted bystander data capture. In care settings also, ethical alignments dictate bolstering safeguards for vulnerable groups and avoiding approaches that make adopting such technologies a precondition of assistance (e.g., coerced surveillance adopted in exchange for care).

4.5 Interoperability, continuity and lock-in: the hidden system risk

Memory infrastructures are durable. If an individual or organisation depends on a platform, switching costs can become exorbitant, leading to lock-in. Interoperability and data portability are not just engineering questions; they are conditions of autonomy and equity. For clinical prosthetics and AI assistants, continuity planning (backups, export formats, migration pathways) becomes part of responsible implementation. Without it, benefits become brittle and failure modes severe (loss of external memory, disruption of care routines).

4.6 Domain implications - technology–society lens

4.6.1 Work and organizations

In workplaces, memory-support tools enable more distributed task management and hand-off; equally, they expose users to greater surveillance and performance pressure. AI summarization of meetings and messages may change the organisation's memories: who gets credit, what is chronicled, what is indexed and searchable. Good governance will clarify the scope of permitted use, time limits for retention, and contestability pathways, particularly where records may be made use of in disciplinary processes.

4.6.2 Education

In education, 'offloading' will substantially change students' learning strategies and study habits. The critical question is not whether students use offloading tools for memory, but how to design systems so that learning outcomes are as optimal as possible whilst managing attention and the constraints that afford deep encoding of material. Policy must draw a line between offloading for logistics, and offloading that produces inadequate coverage of material. Transparency in the extent to which AI supports summarization or citation is needed so that dependence is not hidden.

4.6.3 Healthcare and care-at-home

In healthcare, memory support may improve adherence and continuity; AI-enabled memory assistants holding sensitive personal data will require clinical governance, professional training, and clarity of responsibility boundaries. Adoption should be embedded into established care pathways (selection, training, monitoring, caregiver support).

Privacy-by-design is critical to this space; given the high sensitivity, high consequence of health/life trace data.

4.7 Twelve principles for responsible AI-mediated memory support, implementation-ready

Building on our synthesis, we suggest a dozen implementable principles for AI-mediated memory support, each of which can be understood as being translated into implementation requirements and governance controls.

Designers/developers:

1. Memory support is a workflow, not a feature — minimize friction for legitimate work, maximize prevention of harmful proactivity.
2. Integrate provenance and correction capabilities from the beginning — avoid relying on post-hoc rationalization.
3. Capture learning and attention effects in addition to successful retrieval.
4. Support for accessibility and multiple languages should be built into products by design.

Healthcare professionals/caregivers:

1. Memory technologies are most effective as a service delivery model (e.g., selection, training, monitoring) rather than individual product offerings.
2. In addition to assessing changes in cognition, evaluate caregiver burden and compliance/adherence rates when using these technologies.

Workplace/school/public service organizations:

1. Establish governing roles and guidelines for permissible use prior to deploying memory support technologies — specifically to prevent covert surveillance usage.
2. Where possible, provide training and opt-out options for all users, particularly those who may be more susceptible to coercion.
3. Develop mechanisms for accountability and redress for users so they can dispute inaccurate or misleading records or summaries created through use of memory support technologies.

Regulatory/policy makers:

1. Support development of open, interoperable standards for user-controlled portable personal data formats to enhance user control and independence.
2. Support and fund regulatory and oversight efforts to require adequate audits and transparency relative to perceived risks — and with particular protections related to life-logging and health-related applications.
3. Support funding of independent evaluations and digital literacy initiatives to ensure that the benefits of memory support technologies are not solely dependent upon access to private capital.

4.8 Align implementation with standards and regulation, practical mapping

Given the very real need for some platforms to be working with fragile data and the high potential for harm around it, AI-mediated memory support increasingly requires alignment with established standards and regulatory frameworks, given the sensitivity of personal data and the potential for harm. While detailed legal analysis of obligations is jurisdiction-specific, there are some widely-used instruments whose expectations converge in ways that lend themselves to translation into operational requirements.

NIST AI RMF 1.0 emphasizes risk management across the AI lifecycle, including governance roles, how to measure, and monitoring; their Generative AI Profile includes generative AI systems and reinforces documentation, evaluation, and incident handling as foundational controls (National Institute of Standards and Technology, 2023, 2024).

OECD AI Recommendation lays out principles of inclusive growth, human centred values, transparency, robustness, and accountability, encouraging actual mechanisms for implementation rather than aspirational principles for purely aspirational statements without operational mechanisms (Organisation for Economic Co-operation and Development, 2019).

UNESCO AI Ethics Recommendation focuses on human rights, proportionality, transparency, responsibility, with a focus on capacity-building and oversight as core areas of accountability for vulnerable populations (UNESCO, 2021).

EU AI Act takes a risk-based approach with obligations that scale for higher-risk systems, and strengthens requirements around governance, documentation, transparency, and oversight (European Parliament and Council of the European Union, 2024).

Table 4 illustrates how the system conditions outlined in Table 2 are translated into tangible, verifiable, and deployable implementation components that can provide a basis for in place readiness assessments at the time of activation of AI systems in organizations (e.g. university, hospital, workplace). Rather than treat "readiness" as a declarative statement of compliance; Table 4 outlines specific documentation and process outputs that reviewers can inspect - i.e., RACI matrices that identify ownership and accountability (for example), Change Control Records, Incident Response Runbooks, Data Inventory and Retention Schedules, Integration Diagrams, Backup and Disaster Recovery Plans, Pre-Deployment Test Reports, Usability Trial Summaries, Bias and Subgroup Performance Reports, Post-Deployment Follow-Up Plans. Across the six areas of interest, Security and Reliability are most effectively viewed as cross-cutting properties that are demonstrated through many of the same types of artifacts (i.e., Access Control Policies, Incident Handling Procedures, Resilience Planning). In many cases, the required artifacts will have been previously developed and audited by mature IT and Clinical/Enterprise Risk Functions, making auditing much simpler. Governance and Accountability, Contestability and Redress, Ethical/Legal Alignment typically require some level of role assignments, Decision Rights, Escalation Pathways, End-to-End Operational Workflows, and in many instances, will have been documented but require organizational buy-in and adherence to be effective, thus can reveal gaps that are not strictly technical. Risk Management Frameworks (such as the NIST AI Risk Management Framework (AI RMF 1.0)) provide structure for the identification, measurement, and management of AI related risks including potential legal, ethical and social impacts. High Level Frameworks cannot replace system specific assurance activities and therefore the Evidence Set outlined in Table 4 provides a combination of Governance Artifacts with Technical/Human Factors Evaluation Outputs providing reviewers with a means of assessing whether policies exist and if so whether the system has been tested, monitored and operationalized in a manner that supports the use of the system safely, reliably and equitably.

Table 4. Mapping system-level conditions to concrete implementation artifacts

Condition	Practical artifacts that reviewers/auditors can inspect
Governance & accountability	Named owner and accountability RACI; audit plan; incident response runbook; change-control log.
Privacy-by-design	Data inventory; consent scripts; retention schedule; access control policy; deletion and portability procedure.
Interoperability and infrastructure	Export formats; backup and disaster recovery plan; integration diagram; uptime and recovery targets.
Evaluation and assurance	Pre-deployment test report; usability trial summary; bias and subgroup performance report; longitudinal follow-up plan.
Capacity-building	Training materials; onboarding checklist; support ticket workflow; clinician/caregiver guidance.
Ethical–legal alignment	Risk classification memo; safeguards for vulnerable users; contestability and redress policy; compliance mapping checklist.

4.9 Stakeholder-specific recommendations

To increase translation into practice, we summarize what “responsible memory support” means for key stakeholders.

Designers and developers:

1. Treat memory support as a workflow, not a feature: reduce friction for legitimate tasks while preventing harmful proactivity.
2. Implement provenance and correction mechanisms early; do not rely on post-hoc explanations.
3. Measure attention and learning impacts, not only retrieval success.
4. Build accessibility and multilingual support as first-class requirements.

Clinicians and care providers:

1. Deploy memory technologies as part of a service model (selection, training, monitoring), not as standalone tools. Prefer tools that are portable and continuous so that when care settings change there is less disruption.

Organizations (workplaces, schools, public services):

1. Governance roles and permissible-use policy should be established prior to deployment; uses for surveillance should be prohibited.
2. Provide training and opt-out pathways where appropriate, particularly for vulnerable users.
3. Provide contestability and redress pathways so that individuals can challenge and rectify incorrect records or summaries.

Policymakers and regulators:

1. Secure interoperable, portable personal data formats to avoid lock-in and support individual autonomy.
2. Support auditing and transparency requirements commensurate with risk, especially where life-logging and health are concerned.
3. Fund independent evaluation and digital literacy programmes, so that there is no necessity for the benefits to accrue exclusively to private organisations.

The recommendations above are intended to sit alongside the system level conditions in Table 2, as well as the implementation artefacts in Table 4; when packaged together these are things which may help to reduce abandonment and predictably uneven use and improve legitimate use in real deployments.

4.10 Population-level cognition and life-course exposure to digital memory infrastructures

A perennial question in technology-and-society debates is whether persistent internet access and pervasive external memory infrastructures created a population-level change in cognition. A broad synthesis frames this as the internet potentially changing cognitive habits, particularly around attention and memory, making calls for life-course research that distinguishes transient habit-convenience effect from lasting developmental effects (Firth et al., 2019). For memory-support technologies, this suggests two scopes for evaluation: short-term functional performance (e.g., does it reduce everyday memory failures?) and long-term cognitive ecology (e.g., does sustained offloading change which skills are developed and how, over years?). For IJITSS, this question is a social question: life-course effects may be distributed unfairly. Students, precarious workers and older adults may experience different benefits and harms, because their exposure to tools, their alternatives and their sources of support are different. This adds weight to the case for using equitable use and allocating public investment for digital literacy and accessible support services, rather than hoping they spread widely of their own accord.

5. Limitations

We reviewed a wide variety of evidence types that are not directly comparable; therefore, we cannot provide an overall pooled effect size estimate. Internal validity is supported by controlled experiments; however, generalizability to everyday situations may be limited due to their structured nature. Clinical studies typically involve smaller sample sizes and are highly context-dependent, while governance documents vary in their normative grounding but have practical implications for policymakers. There is limited longitudinal evidence on how long-term cognitive effects of ubiquitous cognitive offloading develop over the life course, and observed relationships may be influenced by factors such as education, occupation, and baseline digital literacy. The limitations identified above inform the proposed research agenda.

6. Research Agenda

We have identified six high-priority areas for additional research that can enhance the quality of our review as well as the applicability of our findings to real-world contexts.

I. Cognition and Self-Efficacy Over Time: Longitudinal, multi-year studies should be conducted to measure the effects of cognitive offloading on the development of encoding strategies, the accuracy of metacognitive calibration, the allocation of attention, and perceptions of self-efficacy within and across cohorts.

II. Effectiveness of Cognitive Offloading Technology in Real-World Contexts: We will investigate the effects of cognitive offloading technologies on real-world behaviors such as task completion, adherence to use guidelines, well-being, and caregiver burden.

III. Equity and Differential Effects: We will examine how socioeconomic status, disability, language proficiency, digital literacy, and prior experiences with technology influence both the positive and negative effects of cognitive offloading and AI-mediated memory support.

IV. Privacy, Informed Consent, and Governance Models for Shared Spaces in Lifelogging: We will develop processes for obtaining informed consent, methods for redacting personally identifiable information from lifelogs, and models for governing shared lifelogging spaces (including the rights of bystanders).

V. Assessing and Auditing the Provenance of AI Memory Pipelines: We will establish auditable processes for measuring the provenance of AI memory pipelines, document best practices for auditing these pipelines, and evaluate the potential for auditing to reduce harm and increase user trust.

VI. Clinical Trials and Service Models: We will compare the performance of AI assistants with that of structured, non-AI cognitive prosthetic services (such as MSS) using criteria including cost, abandonment rates, and caregiver outcomes.

In addition to the six priority areas above, we also recommend conducting mixed-methods studies of cognitive offloading technologies (quantitative outcome measurements supplemented by qualitative data from user and caregiver interviews), as well as sector-specific evaluations of governance (e.g., workplace policies, educational standards, clinical protocols) to better understand the power dynamics and legitimacy of cognitive offloading technologies in a variety of real-world contexts.

7. Conclusions

Memory support technologies have affected cognition and society through the redistribution of “remembering” among individuals, machines, and organizations.

Evidence suggests that memory aids can increase task accomplishment, reduce cognitive load, and foster independence, provided that memory support is accessible (usable), relevant (to tasks or environments), and embedded in regular routines with adequate training and ongoing support. However, these advantages are accompanied by disadvantages. For example, encoding strategies and metacognitive processes may change, users may incur attentional costs due to reliance on the aid, and privacy risks may arise from collecting and storing personal data. In addition, there is the potential for governance failures in deploying these memory aids, especially when AI-based memory aids or lifelogging systems are involved.

The current AI inflection point may substantially increase both the benefits and risks of memory support technology. AI systems can assist with semantic retrieval from large personal archives and generate summaries, but responsible deployment requires clearly defined, transparent, and auditable pipelines, privacy-by-design safeguards, and mechanisms for contestability and correction to prevent opaque harms.

This review presents an evidence-based taxonomy and comparative evaluation matrix to support a practical framework for developing and implementing ethically sustainable memory support in the era of AI.

Declarations

The ethics approval is not required as this review is literature-based and did not involve any human or animal subjects. Funding for this study was not provided by an external source. No conflicts of interest to declare. Data from this study will not be available due to it being a review of previously published studies.

REFERENCES

1. Allé, M. C., Manning, L., Potheegadoo, J., Coutelle, R., Danion, J.-M., & Berna, F. (2017). Wearable cameras are useful tools to investigate and remediate autobiographical memory impairment: A systematic PRISMA review. *Neuropsychology Review, 27*(1), 81–99. https://doi.org/10.1007/s11065-016-9337-x
2. Boman, I.-L., Tham, K., Granqvist, A., Bartfai, A., & Hemmingsson, H. (2007). Using electronic aids to daily living after acquired brain injury: A study of learning and usability. *Disability and Rehabilitation: Assistive Technology, 2*(1), 23–33. https://doi.org/10.1080/17483100600856213
3. Buolamwini, J., & Gebru, T. (2018). Gender shades: Intersectional accuracy disparities in commercial gender classification. *Proceedings of Machine Learning Research, 81*, 77–91. https://proceedings.mlr.press/v81/buolamwini18a.html
4. Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis, 58*(1), 7–19. https://doi.org/10.1093/analys/58.1.7
5. de Fine Licht, K., & de Fine Licht, J. (2020). Artificial intelligence, transparency, and public decision-making: Why explanations are key when trying to produce perceived legitimacy. *AI & Society, 35*(4), 917–926. https://doi.org/10.1007/s00146-020-00960-w
6. Diakopoulos, N. (2015). Algorithmic accountability. *Digital Journalism, 3*(3), 398–415. https://doi.org/10.1080/21670811.2014.976411

7. European Parliament and Council of the European Union. (2024). *Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union.* https://eur-lex.europa.eu/eli/reg/2024/1689/oj
8. Felzmann, H., Villaronga, E. F., Lutz, C., & Tamò-Larrieux, A. (2019). Transparency you can trust: Transparency requirements for artificial intelligence between legal norms and contextual concerns. *Big Data & Society, 6*(1), Article 205395171986054. https://doi.org/10.1177/2053951719860542
9. Firth, J., Torous, J., Stubbs, B., Firth, J. A., Steiner, G. Z., Smith, L., Alvarez-Jimenez, M., Gleeson, J., Vancampfort, D., Armitage, C. J., & Sarris, J. (2019). The "online brain": How the internet may be changing our cognition. *World Psychiatry, 18*(2), 119–129. https://doi.org/10.1002/wps.20617
10. Gebru, T., Morgenstern, J., Vecchione, B., Wortman Vaughan, J., Wallach, H., Daumé, H., III, & Crawford, K. (2018). *Datasheets for datasets* (arXiv:1803.09010). arXiv. https://doi.org/10.48550/arXiv.1803.09010
11. Greenaway, M. C., Duncan, N. L., & Smith, G. E. (2013). The memory support system for mild cognitive impairment: Randomized trial of a cognitive rehabilitation intervention. *International Journal of Geriatric Psychiatry, 28*(4), 402–409. https://doi.org/10.1002/gps.3838
12. Henkel, L. A. (2014). Point-and-shoot memories: The influence of taking photos on memory for a museum tour. *Psychological Science, 25*(2), 396–402. https://doi.org/10.1177/0956797613504438
13. Hodges, S., Williams, L., Berry, E., Izadi, S., Srinivasan, J., Butler, A., Smyth, G., Kapur, N., & Wood, K. (2006). SenseCam: A retrospective memory aid. In P. Dourish & A. Friday (Eds.), *UbiComp 2006: Ubiquitous computing* (Lecture Notes in Computer Science, Vol. 4206, pp. 177–193). Springer. https://doi.org/10.1007/11853565_11
14. Jamieson, M., Cullen, B., McGee-Lennon, M., Brewster, S., & Evans, J. J. (2014). The efficacy of cognitive prosthetic technology for people with memory impairments: A systematic review and meta-analysis. *Neuropsychological Rehabilitation, 24*(3–4), 419–444. https://doi.org/10.1080/09602011.2013.825632
15. Jamieson, M., Cullen, B., McGee-Lennon, M., Brewster, S., & Evans, J. J. (2017). Technological memory aid use by people with acquired brain injury. *Neuropsychological Rehabilitation, 27*(6), 919–936. https://doi.org/10.1080/09602011.2015.1103760
16. Kalnikaitė, V., & Whittaker, S. (2007). Software or wetware? Discovering when and why people use digital prosthetic memory. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 71–80). Association for Computing Machinery. https://doi.org/10.1145/1240624.1240635
17. Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I. D., & Gebru, T. (2019). Model cards for model reporting. In *Proceedings of the Conference on Fairness, Accountability, and Transparency* (pp. 220–229). Association for Computing Machinery. https://doi.org/10.1145/3287560.3287596
18. National Institute of Standards and Technology. (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)* (NIST AI 100-1). https://doi.org/10.6028/NIST.AI.100-1
19. National Institute of Standards and Technology. (2024). *Artificial Intelligence Risk Management Framework: Generative AI Profile* (NIST AI 600-1). https://doi.org/10.6028/NIST.AI.600-1
20. Nissenbaum, H. (2009). *Privacy in context: Technology, policy, and the integrity of social life.* Stanford University Press. https://doi.org/10.1515/9780804772891
21. Organisation for Economic Co-operation and Development. (2019). *Recommendation of the Council on Artificial Intelligence* (OECD/LEGAL/0449). https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449
22. Pauly-Takacs, K., Moulin, C. J. A., & Estlin, E. J. (2011). SenseCam as a rehabilitation tool in a child with severe memory impairment. *Memory, 19*(7), 707–718. https://doi.org/10.1080/09658211.2010.494046
23. Raji, I. D., Smart, A., White, R. N., Mitchell, M., Gebru, T., Hutchinson, B., Smith-Loud, J., Theron, D., & Barnes, P. (2020). Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (pp. 33–44). Association for Computing Machinery. https://doi.org/10.1145/3351095.3372873
24. 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
25. Sachdeva, C., & Gilbert, S. J. (2020). Excessive use of reminders: Metacognition and effort-minimisation in cognitive offloading. *Consciousness and Cognition, 85*, Article 103024. https://doi.org/10.1016/j.concog.2020.103024
26. Scarampi, C., & Gilbert, S. J. (2020). The effect of recent reminder setting on subsequent strategy and performance in a prospective memory task. *Memory, 28*(5), 677–691. https://doi.org/10.1080/09658211.2020.1764974
27. Sellen, A. J., Fogg, A., Aitken, M., Hodges, S., Rother, C., & Wood, K. (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* (pp. 81–90). Association for Computing Machinery. https://doi.org/10.1145/1240624.1240636
28. Sellen, A. J., & Whittaker, S. (2010). Beyond total capture: A constructive critique of lifelogging. *Communications of the ACM, 53*(5), 70–77. https://doi.org/10.1145/1735223.1735243

29. Silva, A. R., Pinho, M. S., Macedo, L., Moulin, C., Caldeira, S., & Firmino, H. (2017). It is not only memory: Effects of SenseCam on improving well-being in patients with mild Alzheimer disease. *International Psychogeriatrics, 29*(5), 741–754. https://doi.org/10.1017/S104161021600243X
30. Silva, A. R., Pinho, M. S., Macedo, L., & Moulin, C. J. A. (2018). A critical review of the effects of wearable cameras on memory. *Neuropsychological Rehabilitation, 28*(1), 117–141. https://doi.org/10.1080/09602011.2015.1128450
31. Sparrow, B., Liu, J., & Wegner, D. M. (2011). Google effects on memory: Cognitive consequences of having information at our fingertips. *Science, 333*(6043), 776–778. https://doi.org/10.1126/science.1207745
32. Storm, B. C., & Stone, S. M. (2015). Saving-enhanced memory: The benefits of saving on the learning and remembering of new information. *Psychological Science, 26*(2), 182–188. https://doi.org/10.1177/0956797614559285
33. UNESCO. (2021). *Recommendation on the Ethics of Artificial Intelligence.* https://www.unesco.org/en/legal-affairs/recommendation-ethics-artificial-intelligence
34. Ward, A. F., Duke, K., Gneezy, A., & Bos, M. W. (2017). Brain drain: The mere presence of one's own smartphone reduces available cognitive capacity. *Journal of the Association for Consumer Research, 2*(2), 140–154. https://doi.org/10.1086/691462