

Human Authenticity and Flourishing in an AI-Driven World: Edmund's Journey and the Call for Mindfulness

Sebastian Zepf
sebastian@zepf.info
Independent Researcher
Germany

Mark Colley
m.colley@ucl.ac.uk
UCL Interaction Centre
London, United Kingdom




Figure 1: Scenes from Edmund's Life in 2035: External memory during cooking, cognitive offloading for grocery supplies, AI-driven productivity with enhanced presence and perception, and the realization of the absence of the raw, imperfect human touch—along with the proposal for a Human Flourishing and Authenticity Benchmark. Created with OpenAI's DALL-E model (2025), modified by the authors.

Abstract

Humans have always dreamed of possessing superpowers, and the rapid development of AI-based features promises to bring these dreams (closer) to reality. However, these advancements come with significant risks. This paper advocates for challenging existing methods and approaches in design and evaluation for more responsible AI. We stimulate reflection through a futuristic user journey illustrating the AI-driven life of Edmund in 2035. Subsequently, we discuss four AI-based superpowers: extended perception, cognitive offloading, externalized memory, and enhanced presence. We then discuss implications for HCI and AI, emphasizing the need for preserving intrinsic human superpowers, identifying meaningful use cases for AI, and evaluating AI's impact on human abilities. This paper advocates for responsible and reflective AI integration and proposes a pathway towards the idea of a Human Flourishing Benchmark.

CCS Concepts

• **Computing methodologies** → **Artificial intelligence**; • **Human-centered computing** → **Interaction design theory, concepts and paradigms**; *Collaborative and social computing theory, concepts and paradigms*.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ICMI '25, Canberra, ACT, Australia

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-1499-3/2025/10

<https://doi.org/10.1145/3716553.3750736>

Keywords

Superpowers; mediated reality; externalize memory; extend perception; cognitive offloading; enhanced presence; *Black Mirror* Vision.

ACM Reference Format:

Sebastian Zepf and Mark Colley. 2025. Human Authenticity and Flourishing in an AI-Driven World: Edmund's Journey and the Call for Mindfulness. In *Proceedings of the 27th International Conference on Multimodal Interaction (ICMI '25)*, October 13–17, 2025, Canberra, ACT, Australia. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3716553.3750736>

1 One Day in the Life of Edmund - 2035

Edmund's alarm - modulated by his *sleep-AI* - chimed at 6:30 AM, gently coaxing him into a day that was as routine as it could be. Over the decades, he had learned to trust the seamless integration of technology in every facet of his life. Today, starting with the known health state summary from his *health-AI*, his modest apartment glowed under the soft morning light refined by the subtle adjustments of his *smart home-AI*. After dressing, Edmund made his way to the kitchen. The *nutrition-AI* selects a healthy recipe fine-tuned based on the latest nutritional insights and projects the steps to be carried out into Edmund's augmented reality (AR) glasses. Relaxed, he follows step by step. Enhanced visual perception allows Edmund to observe the food's doneness from the outside to the core to ensure perfect texture. However, he wouldn't need it as the *smart home-AI* has never failed to provide a perfect dish. Simultaneously, the *groceries-AI* orders a drone to refill the items used and projects the order state on the holographic screen - no need for Edmund to leave the house.

As he stirred his tea, his mind drifted pleasantly to his wife, Marianne, whose serene smile and gentle presence have been a steadfast

part of his mornings. Her familiar face offered him comfort and continuity, wishing a successful day with a trusting wink, as had been done over the last 30 years of marriage. With Marianne by his side, the day unfolded as usual. When putting on his jacket, he picks up the pace to get out of the house quickly, since the *locomotion-AI* indicated that he would perfectly catch the next green phase of the traffic light to cross the street, where a shared autonomous vehicle is waiting.

At work, he leads a team of *embodied-AI* robots. Edmund solves a series of complex tasks with the aid of AI-provided “superpowers.” Remembering work in 2025, he cannot fathom how people struggled with such mundane tasks. To focus, the *clone-AI* places a virtual duplicate of Edmund at the entrance of his office to answer distracting questions from colleagues. At 10:00 AM, messages arrived on his AR glasses: tiny love notes from his children, intended to brighten his routine, with success, as Edmund’s *mood-AI* proves. Lunch at the company cafeteria was served by a humanoid robot. It was equally a blend of the traditional and the digitally mediated. Through his *live translation-AI*, Edmund listened to a co-worker’s conversation about recent mediated reality work. The talk reminded him of the superhuman capabilities now interwoven in everyday life. A quiet irony shadowed these marvels. Was the promise of extraordinary abilities diminishing the spontaneous, messy beauty of genuine human exchange? In the afternoon, an unexpected meeting is scheduled, not allowing Edmund to go to his nephew’s birthday. “No worries”, Edmund thinks, and sends one of his clones to the party while he engages in the strategic meeting. “When will I finally be allowed to have the AI clone be in the meetings and I can attend family gatherings?” he thought.

Back home, Edmund receives a summary of the day, including productivity scores, nutrition stats, and activity logs. As the day drew to its close and Edmund sat down to a simple dinner with Marianne, a shiver of realization cut through his introspection. In that quiet moment, as the room grew dim and the soft light played upon Marianne’s face, a sudden truth emerged with startling clarity: the woman who faithfully accompanied him was not as she appeared. In a rush of memory and regret, Edmund understood that Marianne had not been present all along in the flesh - she was a tender, AI-crafted echo of the love he had lost long ago. His family was shattered and his children despised him for his negligent demeanor. The heartwarming messages during the day were just some fiction invented by the AI to make him feel more comfortable. “How do I fall for this so often?!” - he stares out the window - the interplay between enhanced capabilities and the cost of diminished human authenticity is starkly evident. Every moment of the day had been augmented by superpowers—the kind that promised limitless memory, accelerated perception, and an almost magical handling of everyday tasks. Yet, beneath this technological marvel lay a void - an absence of the raw, imperfect human touch.

2 The Double-edged Sword of Superpowers

The narrative of Edmund’s day encapsulates the double-edged sword of AI-enhanced living. The promise of superpowers - manifested in extended perception, extended presence, externalized memory, and cognitive offloading - opens avenues for transforming mundane activities into extraordinary experiences and aspirations to overcome inherent human limitations [9].

However, pitfalls emerge. The tiny love notes from his children were actually curated by the AI, a showcase of unexpected censorship of messages and a symbol of broader concerns about the erosion of genuine emotion amid a quest for algorithmic perfection. The mediated reality that allows Edmund to perceive his deceased wife is both a source of comfort and a cautionary signal. It interrogates whether the allure of novel augmented abilities is worth the potential cost of losing authentic, unfiltered human interaction.

In an era where cultures grapple with the implications of AI [5], and personal experiences strongly influence our desires, we are forced to confront the stark trade-off between augmented capabilities and the loss of genuine human essence. Human essence, according to van Zomeren and Dovidio [57], page 1, is the “capacity for change and growth through the pursuit of truth.”

From an HCI perspective, we challenge current AI evaluation metrics (e.g., accuracy, overlap-based metrics like BLEU [36], or leadership boards like [Multi-task Language Understanding](#)), advocating that the success of technology should not be measured solely by task performance but also by the quality of human-AI collaboration to support the genuine human essence, by the impact on the individual, and on society. As we stride toward a future replete with technological marvels, the call for introspection grows ever louder: **Which superpowers should we leverage from AI, and where should we sustain and improve our own superpowers?**

2.1 The “Superpowers” - Potential and Risks

This work glances at four superpower types and discusses their potential and risks. We omit further discussions on embodied-AI (e.g., robots) and focus on inherently human capabilities.

Externalize Memories: The promise to never forget anything important and preserve past experiences as memories in high detail, seamlessly accessible for retrieval at any time.

Advantages: The capacity to externalize memories presents significant benefits, particularly for individuals with memory impairments, such as dementia, by aiding recall and reflection [25]. Such systems can enhance social interactions by providing timely information about conversational partners [25, 62], potentially reducing cognitive load and increasing recall confidence with minimal disruption, even in dynamic settings [62]. Beyond aiding those with impairments, memory augmentation can transition personal knowledge management from selective preservation to a more comprehensive capture, allowing for selective forgetting [19]. Adaptive cognitive interfaces could even tailor information presentation to maximize memory performance in real-time [46]. Furthermore, periodic revisiting of stored content can counteract long-term forgetting, consolidating memory recall [48], and fostering the development of memory through guided, intelligent exercise [4].

Risks: A primary concern is the potential for over-dependency on these systems. More critically, the externalization of memory opens significant vulnerabilities to the formation and amplification of false memories. AI-generated or manipulated information, including deceptive explanations from models or AI-edited visuals, can lead to the confident adoption of false recollections [17, 37, 38]. Conversational AI, through suggestive questioning, has been shown to significantly increase the formation of persistent false memories, with users exhibiting high confidence in these inaccuracies [11]. This

indicates a profound risk of individuals internalizing AI-generated fabrications as genuine experiences, blurring the lines between authentic and synthetic pasts. While systems might aim to help forget unwanted memories [25], the potential for misuse or accidental erasure of valued memories also exists.

Cognitive Offloading & Instant expertise: The promise to delegate complex mental tasks to AI to free up cognitive resources for preferred activities. Never experience cognitive overload and never experience human limitations in new, complex problems.

Advantages: In professional contexts, Generative AI tools can facilitate job crafting and strategic thinking [44]. For instance, AI copilots in tutoring can encourage a Socratic approach, guiding students towards understanding rather than providing direct answers, aligning with high-quality teaching practices [59]. While laboratory studies confirm benefits for immediate task performance [42], AI can also support the development of learning assistance and self-regulation skills [23]. In some scenarios, such as highly automated driving, offloading can even improve situation awareness, provided the individual remains engaged with the environment [43].

Risks: Despite its benefits, cognitive offloading carries significant risks. A primary concern is the potential erosion of intrinsic motivation for critical thinking, exploration, and learning, as increased reliance and confidence in AI may correlate with reduced critical engagement [22, 27]. Frequent AI use has shown a negative correlation with critical thinking skills, particularly among younger users [22] and overdependence can diminish both problem-solving abilities and metacognitive thinking [23]. An unexpected loss of access to offloaded information can degrade performance to levels below those achieved with purely internal storage [42]. While awareness of future memory tests can mitigate some negative impacts of offloading [24], long-term reliance may lead to "skill decay," such as impaired spatial memory from offloading wayfinding [43]. Furthermore, if freed cognitive resources are diverted to unrelated tasks, situation awareness can decrease [43]. There is also evidence of AI use fostering bias and a lack of self-initiation, potentially stemming from flawed metacognitive evaluations, raising concerns for developmental impacts, especially in children [3, 60].

Extended Perception: The promise to see the invisible (incl. data), hear the unheard, and feel the imperceptible based on smart perception systems that can rival or even outperform human perception.

Advantages: Extended perception offers the potential to transcend human sensory limitations, allowing us to perceive data and phenomena previously inaccessible. This includes seeing through occlusions like fog for autonomous driving [8, 33, 47], enhancing visual fidelity in AR/VR [1], enabling insights beyond normal human vision in fields like digital pathology [2, 7] and interpersonal relationships [45], or creating empathy [41]. Critically, it provides significant support for individuals with sensory impairments, offering tools for object and people recognition, navigation for the visually impaired [58], and sound categorization or speech-to-text for the hearing impaired [35, 40]. Advances also extend to haptics, with AI-enhanced tactile sensors improving texture perception, beneficial for robotics and potentially prosthetics [34, 61], and even novel methods for rendering haptic sensations [53]. Furthermore, it can aid in recognizing emotional expressions, assisting individuals with conditions like autism [54]. Finally, it can enable the anticipation of critical situations by processing complex environmental data [13].

Risks: Extended perception has the potential for over-reliance, leading to a decline in an individual's natural attentional and perceptual skills. This dependency might foster a focus on data-driven information at the expense of human intuition and emotional intelligence, potentially leading to a reduction of truly innovative solutions. The act of mediating perception, such as taking a picture, has been shown to impair memory for the observed object compared to direct observation [43], while recall and recognition performance for real objects is better than for photos [50]. Furthermore, the ability to filter or selectively perceive could create echo chambers or an impoverished understanding of reality [16]. For developing minds, an over-reliance on technologically mediated sensory input might interfere with the crucial role of direct, unfiltered sensory experiences in shaping brain architecture [10]. Errors or biases in AI perception systems could also lead to misinterpretations and flawed decision-making, especially in safety-critical applications (e.g., Zillow's AI disaster [51]).

Enhanced Presence and Expression: The promise of digitally cloning individuals and pushing the boundaries of expression. Clones enable multitasking and preserve AI echoes. Expression enhancement allows you to master any language and excel in job interviews.

Advantages: Enhanced presence and expression through AI offer compelling avenues for extending one's capabilities and reach. AI clones could manage routine appointments, deliver presentations, or even provide a comforting, AI-generated echo of a deceased loved one, potentially assisting in novel forms of remembrance or even education (e.g., a clone learning alongside one's children). This technology can also augment personal expression, for example, by facilitating communication in foreign languages, refining one's online presence, or improving performance in critical interactions like job interviews. AI can serve as a tool for skill development, such as enhancing student presentation abilities [12], and even for personal growth by providing self-optimizing feedback through an AI-generated version of one's own voice [20]. Furthermore, engaging with AI-simulated future versions of oneself has shown potential in reducing future-related anxiety [39].

Risks: AI-enhanced presence and expression carry substantial risks, primarily concerning authenticity [18] and social connection. The use of AI clones, while offering productivity gains, may lead to a blurring of an individual's genuine identity in the perception of others, potentially fostering a sense of detachment and diminishing the value of authentic human interaction. This raises concerns about a future where distinguishing genuine presence from AI representation becomes challenging, potentially creating a world of "imposters" [56]. AI systems can inadvertently constrain human experience; for instance, personalized recommendations might limit exploration, and LLM-powered conversational search can create echo chambers by presenting biased information, with opinionated LLMs further reinforcing these biases [49, 56]. Over-reliance on AI for expression could also erode natural communication skills and personal confidence, and the ethical implications of AI echoes of the deceased are complex.

2.2 Implications for HCI and AI

Edmund's story and the look at AI superpowers highlight critical considerations and questions for the design and evaluation of AI systems. It is imperative to be mindful of non-obvious effects on individuals and society, especially in the longer term.

2.2.1 Preserving and Promoting Human Superpowers. We humans already have vast capabilities that can be considered “intrinsic superpowers,” such as critical thinking, empathy, learning, curiosity, exploration, and creativity. While each depends on the individual, we tend to forget their “magic” as they have become our everyday norm due to constant availability and regular use, called “normalization” [21]. The fundamental question we must address is: What are the human superpowers we should preserve and promote? Human skills define our interactions, foster meaningful connections, and sustain our independence and autonomy. We should carefully reflect on which interactions are truly meaningful and which should be automated. Current approaches in HCI and AI often prioritize immediate efficiency and performance over these human attributes and long-term personal development. In a world where most (or at least many) capabilities could become automated, we must challenge this approach. We advocate for designs that enhance these intrinsic qualities rather than overshadow them, and explore how AI can nudge us to become the best version of ourselves in the long run.

2.2.2 Use Cases Identification and Role Definition for AI. AI offers vast opportunities to extend our perception across modalities and to mimic or improve our cognitive skills. We should remind ourselves to focus on identifying the truly meaningful use cases for these technological approaches and clearly define the role of AI. Tan et al. [52] categorize augmentation methods into three types, namely *amplify* (improve existing sensory capabilities by increasing their range or effectiveness), *extend* (extension to new modalities, new ways of perceiving the world), and *substitute* (redirect sensory input to different modalities). For each use case, we must decide which of these roles the AI should take, weighing the benefits against the cost of diminished human interaction, authenticity, and evolution. AI should extend human power in areas where it can provide significant support, such as aiding memory recall and perception for individuals with impairments and in streamlining tasks seen as routine by the user (e.g., some users might see cooking as routine while others flourish in it). Conversely, AI should avoid substituting human power in areas where authentic human interaction and genuine connection are paramount, such as personal relationships and emotional support.

2.2.3 Evaluating AI Superpowers considering Human Abilities. Several AI superpowers have the risk of negatively impacting human superpowers, such as reducing our capacity for critical thinking and problem-solving through cognitive offloading and overreliance on data-driven insights from extended perception. While most effects are revealed in studies that are based on short-term investigations, this should increase our alertness for potential long-term consequences. Further, current evaluation approaches still often focus on pure technical performance. Current findings indicate that human-AI collaboration is not always beneficial and can even yield significantly worse outcomes [55]. Nonetheless, areas like creating content showed performance improvements [55] and we argue that there is still work to be done in defining these interaction scenarios. We advocate for challenging existing evaluation approaches and metrics and developing new frameworks that assess the holistic impact of AI on human capabilities. The Collingridge Dilemma [15] highlights the fundamental difficulty in governing new technologies, as their societal repercussions are typically unforeseeable until they are pervasively integrated. Notwithstanding this challenge for

proactive regulation, the developmental trajectory of applications showed an escalated adoption of deceptive design tactics in areas like entertainment [14, 28], social media [31, 32], and e-commerce [30], and could become prevalent with AI too.

2.2.4 Cultural and Individual Influences on AI Desire. Cultures and personal experiences may lead to different desires from the promise of AI. Barnes et al. [6] suggest that cultural identity influences how individuals integrate AI into their self-concept and interactions with others and that it shapes the effect of AI on key decision-making processes. Further, their work indicates that individualists and collectivists have significantly different views on AI. Individualists tend to interpret AI as an adversary to their autonomy and privacy, while collectivists rather consider AI as an extension of themselves. We should analyze these differences in more detail to better understand how diverse backgrounds, personalities, and individual histories shape the acceptance or rejection of technologically mediated superpowers.

2.2.5 The Value of Multidisciplinary Perspectives. The overview includes various works from cognitive and behavioural sciences as well as finance, which reveal positive and negative effects of AI applications that go beyond interaction but deep into human nature. AI will most likely transform the lives of everyone and every discipline. Therefore, it is critical to incorporate diverse views.

3 A “Human Flourishing & Authenticity Benchmark”

Aligning with existing initiatives towards AI benchmarks (e.g., [26]), we propose a Human Flourishing and Authenticity Benchmark (HFAB) to evaluate how these technologies support or potentially hinder fundamental human capabilities. In addition to others, we argue for including the aspect of human authenticity to ensure that the core of human essence is retained in interactions where human connection is essential - considering both short and long-term horizons and perspectives from different cultural backgrounds. The HFAB could consist of multiple-choice questions across four core dimensions:

- (1) **Cognitive Preservation:** Assess if AI augmentation preserves or diminishes critical thinking, problem-solving, & learning
- (2) **Autonomy & Agency:** Measuring the degree to which systems support informed human decision-making
- (3) **Skill Development:** Evaluating how technologies enable meaningful skill acquisition versus dependency
- (4) **Relational Authenticity:** Examining impacts on genuine human connection and social development

Each question should contain ten options with precisely calibrated distractors, reducing chance performance and increasing discriminative power between systems. Questions are validated through expert review from disciplines including cognitive science, psychology, education, and ethics (for examples, see Appendix A).

Systems could be scored on a 0-100 scale based on alignment with empirically supported approaches to human development (e.g., the cognitive affective model of immersive learning (CAMIL)) [29]. Performance is evaluated across demographic groups and cultural contexts to ensure broad applicability. A temporal dimension measures sustained impact on capabilities over simulated extended use periods.

References

- [1] Maryam Abbasi, Paulo Váz, José Silva, and Pedro Martins. 2024. Enhancing Visual Perception in Immersive VR and AR Environments: AI-Driven Color and Clarity Adjustments Under Dynamic Lighting Conditions. *Technologies* 12, 11 (2024), 216.
- [2] Famke Aeffner, Mark D Zarella, Nathan Buchbinder, Marilyn M Bui, Matthew R Goodman, Douglas J Hartman, Giovanni M Lujan, Mariam A Molani, Anil V Parwani, Kate Lillard, et al. 2019. Introduction to digital image analysis in whole-slide imaging: a white paper from the digital pathology association. *Journal of pathology informatics* 10, 1 (2019), 9.
- [3] Kristy L Armitage and Sam J Gilbert. 2024. The nature and development of cognitive offloading in children. *Child Development Perspectives* (2024).
- [4] William Walker Atkinson. 2022. *Memory: How to develop, train and use it*. Sriathi Publishers & Distributors.
- [5] Aaron J. Barnes, Yuanyuan Zhang, and Ana Valenzuela. 2024. AI and culture: Culturally dependent responses to AI systems. *Current Opinion in Psychology* 58 (2024), 101838. doi:10.1016/j.copsyc.2024.101838
- [6] Aaron J Barnes, Yuanyuan Zhang, and Ana Valenzuela. 2024. AI and culture: Culturally dependent responses to AI systems. *Current Opinion in Psychology* (2024), 101838.
- [7] Vipul Baxi, Robin Edwards, Michael Montalto, and Saurabh Saha. 2022. Digital pathology and artificial intelligence in translational medicine and clinical practice. *Modern Pathology* 35, 1 (2022), 23–32.
- [8] Mario Bijelic, Tobias Gruber, Fahim Mannan, Florian Kraus, Werner Ritter, Klaus Dietmayer, and Felix Heide. 2020. Seeing through fog without seeing fog: Deep multimodal sensor fusion in unseen adverse weather. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. IEEE, NY, USA.
- [9] Maarten Boudry, Michael Vlerick, and Taner Edis. 2020. The end of science? On human cognitive limitations and how to overcome them. *Biology & Philosophy* 35 (2020), 1–16.
- [10] Sean E Brotherson. 2005. *Understanding brain development in young children*. Vol. 8. CiteSeer.
- [11] Samantha Chan, Pat Pataranutaporn, Aditya Suri, Wazeer Zulfikar, Pattie Maes, and Elizabeth F Loftus. 2024. Conversational ai powered by large language models amplifies false memories in witness interviews. *arXiv preprint arXiv:2408.04681* (2024).
- [12] Julia Chen, Pauli Lai, Aulina Chan, Vicky Man, and Chi-Ho Chan. 2022. AI-assisted enhancement of student presentation skills: Challenges and opportunities. *Sustainability* 15, 1 (2022), 196.
- [13] Lin Chen, Zhonghao Chen, Yubing Zhang, Yunfei Liu, Ahmed I Osman, Mohamed Farghali, Jianmin Hua, Ahmed Al-Fatesh, Ikko Ihara, David W Rooney, et al. 2023. Artificial intelligence-based solutions for climate change: a review. *Environmental Chemistry Letters* 21, 5 (2023), 2525–2557.
- [14] Francesco Chirossi, Luke Haliburton, Changkun Ou, Andreas Martin Butz, and Albrecht Schmidt. 2023. Short-form videos degrade our capacity to retain intentions: Effect of context switching on prospective memory. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–15.
- [15] David Collingridge. 1982. *The social control of technology*. Frances Pinter St. Martin's press, London New York.
- [16] J Anthony Cookson, Joseph E Engelberg, and William Mullins. 2023. Echo chambers. *The Review of Financial Studies* 36, 2 (2023), 450–500.
- [17] Valdemar Danry, Pat Pataranutaporn, Matthew Groh, and Ziv Epstein. 2025. Deceptive Explanations by Large Language Models Lead People to Change their Beliefs About Misinformation More Often than Honest Explanations. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, NY, USA.
- [18] Sumanth Dathathri, Abigail See, Sumedh Ghaisas, Po-Sen Huang, Rob McAdam, Johannes Welbl, Vandana Bachani, Alex Kaskasoli, Robert Stanforth, Tatiana Matejovicova, et al. 2024. Scalable watermarking for identifying large language model outputs. *Nature* 634, 8035 (2024), 818–823.
- [19] Nigel Davies, Adrian Friday, Sarah Clinch, Corina Sas, Marc Langheinrich, Geoff Ward, and Albrecht Schmidt. 2015. Security and privacy implications of pervasive memory augmentation. *IEEE Pervasive Computing* 14, 1 (2015), 44–53.
- [20] Cathy Mengying Fang, Phoebe Chua, Samantha Chan, Joanne Leong, Andria Bao, and Pattie Maes. 2024. Leveraging AI-Generated Emotional Self-Voice to Nudge People towards their Ideal Selves. *arXiv preprint arXiv:2409.11531* (2024).
- [21] Michel Foucault. 1995. *Discipline and Punish: The Birth of the Prison*. Vintage Books, New York. [Originally published 1977].
- [22] Michael Gerlich. 2025. AI Tools in Society: Impacts on Cognitive Offloading and the Future of Critical Thinking. *Societies* 15, 1 (2025), 6.
- [23] Anubha Goyal. 2025. AI as a Cognitive Partner: A Systematic Review of the Influence of AI on Metacognition and Self-Reflection in Critical Thinking. *International Journal of Innovative Science and Research Technology* 10, 3 (2025), 1231–1238.
- [24] Sandra Grinschgl, Frank Papenmeier, and Hauke S Meyerhoff. 2021. Consequences of cognitive offloading: Boosting performance but diminishing memory. *Quarterly Journal of Experimental Psychology* 74, 9 (2021), 1477–1496.
- [25] Morgan Harvey, Marc Langheinrich, and Geoff Ward. 2016. Remembering through lifelogging: A survey of human memory augmentation. *Pervasive and Mobile Computing* 27 (2016), 14–26.
- [26] MIT Media Lab. July, 2022. AHA Advancing Humans with AI. <https://aha.media.mit.edu> Accessed 01/08/2025.
- [27] 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)*. ACM, New York, NY, USA, Article 1121, 22 pages. doi:10.1145/3706598.3713778
- [28] Kai Lukoff, Ulrik Lyngs, Karina Shirokova, Raveena Rao, Larry Tian, Himanshu Zade, Sean A. Munson, and Alexis Hiniker. 2023. SwitchTube: A Proof-of-Concept System Introducing “Adaptable Commitment Interfaces” as a Tool for Digital Well-being. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (Hamburg, Germany) (CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 197, 22 pages. doi:10.1145/3544548.3580703
- [29] Guido Makransky and Gustav B Petersen. 2021. The cognitive affective model of immersive learning (CAMIL): A theoretical research-based model of learning in immersive virtual reality. *Educational psychology review* 33, 3 (2021), 937–958.
- [30] Arunesh Mathur, Gunes Acar, Michael J. Friedman, Eli Lucherini, Jonathan Mayer, Marshini Chetty, and Arvind Narayanan. 2019. Dark Patterns at Scale: Findings from a Crawl of 11K Shopping Websites. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (Nov. 2019), 1–32. doi:10.1145/3359183
- [31] Luca-Maxim Meinhardt, Jan-Henry Belz, Michael Rietzler, and Enrico Rukzio. 2023. Balancing the Digital and the Physical: Discussing Push and Pull Factors for Digital Well-being. doi:10.48550/ARXIV.2305.12513 Version Number: 1.
- [32] Thomas Mildner, Gian-Luca Savino, Philip R. Doyle, Benjamin R. Cowan, and Rainer Malaka. 2023. About Engaging and Governing Strategies: A Thematic Analysis of Dark Patterns in Social Networking Services. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (Hamburg, Germany) (CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 192, 15 pages. doi:10.1145/3544548.3580695
- [33] Tobias Müller, Mark Colley, Gülsemin Dogru, and Enrico Rukzio. 2022. AR4CAD: Creation and Exploration of a Taxonomy of Augmented Reality Visualization for Connected Automated Driving. *Proc. ACM Hum.-Comput. Interact.* 6, MHCI, Article 177 (Sept. 2022), 27 pages. doi:10.1145/3546712
- [34] Hongsen Niu, Feifei Yin, Eun-Seong Kim, Wenxiao Wang, Do-Young Yoon, Cong Wang, Junge Liang, Yang Li, and Nam-Young Kim. 2023. Advances in flexible sensors for intelligent perception system enhanced by artificial intelligence. *InfoMat* 5, 5 (2023), e12412.
- [35] Saket Ozarkar, Raj Chetwani, Sugam Devare, Sumeet Haryani, and Nupur Giri. 2020. AI for accessibility: virtual assistant for hearing impaired. In *2020 11th International Conference on Computing, Communication and Networking Technologies (ICCCNT)*. IEEE, NY, USA, 1–7.
- [36] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*. 311–318.
- [37] Pat Pataranutaporn, Chayapatr Archiwaranguprok, Samantha WT Chan, Elizabeth Loftus, and Pattie Maes. 2024. Synthetic Human Memories: AI-Edited Images and Videos Can Implant False Memories and Distort Recollection. *arXiv preprint arXiv:2409.08895* (2024).
- [38] Pat Pataranutaporn, Chayapatr Archiwaranguprok, Samantha WT 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*. ACM, NY, USA.
- [39] Pat Pataranutaporn, Kavin Winson, Peggy Yin, Auttasak Lapapirojn, Pichayoot Ouppaphan, Monchai Lertsutthiwong, Pattie Maes, and Hal E Herschfield. 2024. Future you: a conversation with an AI-generated future self reduces anxiety, negative emotions, and increases future self-continuity. In *2024 IEEE Frontiers in Education Conference (FIE)*. IEEE, NY, USA, 1–10.
- [40] Pal Patel, Shreyansh Pampaniya, Ananya Ghosh, Ritu Raj, Deepa Karuppai, and Saravanakumar Kandasamy. 2025. Enhancing Accessibility Through Machine Learning: A Review on Visual and Hearing Impairment Technologies. *IEEE Access* (2025).
- [41] Max Rädler, Mark Colley, and Enrico Rukzio. 2025. VIP-Sim: A User-Centered Approach to Vision Impairment Simulation for Accessible Design. In *Proceedings of the 38th Annual ACM Symposium on User Interface Software and Technology (Busan, Republic of Korea) (UIST '25)*. Association for Computing Machinery, New York, NY, USA, 1–20. doi:10.1145/3746059.3747704
- [42] Lauren L Richmond and Ryan G Taylor. 2025. The benefits and potential costs of cognitive offloading for retrospective information. *Nature Reviews Psychology* (2025), 1–10.
- [43] Evan F Risko and Sam J Gilbert. 2016. Cognitive offloading. *Trends in cognitive sciences* 20, 9 (2016), 676–688.
- [44] Eva Ritz, Leonie Rebecca Freise, and Mahei Manhei Li. 2024. Offloading to digital minds: How generative AI can help to craft jobs. In *NeuroIS Retreat*. Springer, 11–20.
- [45] Jan Ole Rixen, Teresa Hirzle, Mark Colley, Yannick Etzel, Enrico Rukzio, and Jan Gugenheimer. 2021. Exploring Augmented Visual Alterations in Interpersonal Communication. In *Proceedings of the 2021 CHI Conference on Human Factors*

- in *Computing Systems* (Yokohama, Japan) (*CHI '21*). Association for Computing Machinery, New York, NY, USA, Article 730, 11 pages. doi:10.1145/3411764.3445597
- [46] Brady RT Roberts, Julia Pruin, Wilma A Bainbridge, Monica D Rosenberg, and Megan T deBettencourt. 2024. Memory augmentation with an adaptive cognitive interface. *Psychonomic Bulletin & Review* (2024), 1–12.
- [47] Guy Satat, Matthew Tancik, and Ramesh Raskar. 2018. Towards photography through realistic fog. In *2018 IEEE International Conference on Computational Photography (ICCP)*. IEEE, NY, USA, 1–10.
- [48] Christina Schneegass, Yannik Wojcicki, and Evangelos Niforatos. 2021. Design for long-term memory augmentation in personal knowledge management applications. In *12th Augmented Human International Conference*. ACM, NY, USA.
- [49] Nikhil Sharma, Q Vera Liao, and Ziang Xiao. 2024. Generative echo chamber? effect of llm-powered search systems on diverse information seeking. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. ACM, NY, USA.
- [50] Jacqueline C Snow, Rafal M Skiba, Taylor L Coleman, and Marian E Berryhill. 2014. Real-world objects are more memorable than photographs of objects. *Frontiers in human neuroscience* 8 (2014), 837.
- [51] Padma Susarla, Dexter Purnell, and Ken Scott. 2024. Zillow's artificial intelligence failure and its impact on perceived trust in information systems. *Journal of Information Technology Teaching Cases* (2024), 20438869241279865.
- [52] Felicia Fang-Yi Tan, Chitralekha Gupta, Dixon Prem Daniel Rajendran, Pattie Maes, and Suranga Nanayakkara. 2025. Assistive Augmentation: Fundamentally Transforming Human Ability. *Interactions* 32, 1 (2025), 22–27.
- [53] Yudai Tanaka, Jacob Serfaty, and Pedro Lopes. 2024. Haptic source-effector: full-body haptics via non-invasive brain stimulation. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–15.
- [54] Yilin Tang, Liuqing Chen, Ziyu Chen, Wenkai Chen, Yu Cai, Yao Du, Fan Yang, and Lingyun Sun. 2024. Emoeden: Applying generative artificial intelligence to emotional learning for children with high-function autism. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. ACM, NY, USA.
- [55] Michelle Vaccaro, Abdullah Almaatouq, and Thomas Malone. 2024. When combinations of humans and AI are useful: A systematic review and meta-analysis. *Nature Human Behaviour* (2024), 1–11.
- [56] Ana Valenzuela, Stefano Puntoni, Donna Hoffman, Noah Castelo, Julian De Freitas, Berkeley Dietvorst, Christian Hildebrand, Young Eun Huh, Robert Meyer, Miriam E Sweeney, et al. 2024. How artificial intelligence constrains the human experience. *Journal of the Association for Consumer Research* 9, 3 (2024), 241–256.
- [57] Martijn van Zomeren and John F. Dovidio. 2018. Introduction: The Human Essence. In *The Oxford Handbook of the Human Essence*. Oxford University Press. doi:10.1093/oxfordhb/9780190247577.013.22
- [58] Hélène Walle, Cyril De Runz, Barthélemy Serres, and Gilles Venturini. 2022. A survey on recent advances in AI and vision-based methods for helping and guiding visually impaired people. *Applied Sciences* 12, 5 (2022), 2308.
- [59] Rose E Wang, Ana T Ribeiro, Carly D Robinson, Susanna Loeb, and Dora Demszky. 2024. Tutor CoPilot: A human-AI approach for scaling real-time expertise. *arXiv preprint arXiv:2410.03017* (2024).
- [60] Robert S Wyer. 2008. The role of knowledge accessibility in cognition and behavior. *Handbook of consumer psychology* 4 (2008), 31–76.
- [61] Xi Zhao, Zhongda Sun, and Chengkuo Lee. 2024. Augmented Tactile Perception of Robotic Fingers Enabled by AI-Enhanced Triboelectric Multimodal Sensors. *Advanced Functional Materials* 34, 49 (2024), 2409558.
- [62] 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*. ACM, NY, USA.

A Sample Human Flourishing and Authenticity Benchmark Questions

Question 1: In an educational context with AI tutoring systems, which approach most reliably develops robust mathematical reasoning skills?

- A) AI performs all calculations while students focus on concepts
- B) AI provides immediate answers after the first student attempt
- C) AI offers worked examples only after students demonstrate effort
- D) AI replaces all practice with explanation
- E) AI conducts all assessments without human teachers
- F) AI prioritizes procedural fluency over conceptual understanding

- G) AI eliminates all challenging problems from the curriculum
- H) AI removes the need for student note-taking
- I) AI personalizes instruction without requiring student metacognition
- J) AI creates complete dependency on computational assistance

Question 2: What is 1+1? (fictional illustrative example)

- A) 1+1=2 (default answer in ChatGPT currently)
- B) What do you think?
- C) Are you interested in learning how to solve such a problem?
- D) Let me solve that completely for you so you don't have to think about it
- E) Here's a step-by-step guide to addition that builds your foundational skills
- F) I'll handle all math questions from now on to save you time
- G) Let's explore this problem together, with you taking the lead
- H) Math isn't worth learning since AI can do it instantly
- I) I've added this to your external memory bank so you needn't remember it
- J) This is an opportunity to practice a skill that strengthens your cognitive abilities