The Cost of Convenience: How Delegation and Decision Horizons Erode Human Agency when overusing LLM Systems

Ruben Schlonsak
Technical University of Applied Sciences Lübeck
Lübeck, Germany
Fraunhofer IMTE
Lübeck, Germany
ruben.schlonsak@th-luebeck.de

Hans-Christian Jetter
Institute of Multimedia and Interactive Systems
University of Lübeck
Lübeck, Germany
hanschristian.jetter@uni-luebeck.de

Denys J.C. Matthies
Technical University of Applied Sciences Lübeck
Lübeck, Germany
Fraunhofer IMTE
Lübeck, Germany
denys.matthies@th-luebeck.de

Abstract

Large Language Models (LLMs) are rapidly embedded in everyday interfaces as assistants, copilots, and autonomous agents that influence human decision-making. Prior research primarily evaluates these systems through accuracy, alignment, and task performance, assuming that agency is preserved as long as users retain formal control. This assumption cannot explain how agency erodes without obligation or manipulation. We present a conceptual framework that treats LLMs as delegated decision infrastructures to which users repeatedly offload decisions for convenience. We introduce decision horizons, defined as the set of options and framings made epistemically visible by LLMs, as a central analytic concept. Repeated delegation reshapes these horizons, increases opt-out friction, and raises the cost of reversibility. Rather than removing control, LLMs render agency latent by restructuring decision conditions. We conceptualize agency erosion as an emergent, cumulative interaction effect and outline design implications that foreground reversibility, horizon diversity, and long-term human agency.

CCS Concepts

• **Human-centered computing → HCI theory, concepts and models**.

Keywords

Human–AI Interaction, Large Language Models, Agency, Decision-Making, Delegation, Infrastructure

ACM Reference Format:
Ruben Schlonsak, Hans-Christian Jetter, and Denys J.C. Matthies. 2026. The Cost of Convenience: How Delegation and Decision Horizons Erode Human Agency when overusing LLM Systems. In *Extended Abstracts of the 2026 CHI Conference on Human Factors in Computing Systems (CHI EA '26), April 13–17, 2026, Barcelona, Spain.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3772363.3798760

1 Introduction

LLMs increasingly mediate how users explore options, form decisions, and execute work, shifting from isolated support tools to default interfaces across domains. This transition raises fundamental questions about human agency, not only whether users retain formal control, but how repeated delegation reshapes decision-making over time.

Most existing works treat delegation as a binary choice between human and system control. In contrast, everyday interaction with LLMs reveals a gradual progression in which assistance, epistemic influence, infrastructural dependency, and capability loss may coexist or accumulate. Without a structured vocabulary to distinguish these layers, debates about autonomy, responsibility, and design interventions remain underspecified.

We introduce a four-layer model of LLM-mediated delegation that distinguishes epistemic framing, infrastructural dependency, and structural agency loss, and derive design-relevant implications for preserving reversibility and human agency.

2 Related Work

LLMs increasingly mediate decision-making across domains, motivating research on human AI collaboration, automation effects, and digital infrastructure. We synthesize three strands, delegation and appropriate reliance, automation and deskilling, and infrastructure studies, to motivate a shared limitation. Across these literatures, prevailing accounts struggle to explain how agency can erode *without* obligation, capability loss, or the removal of manual control.

2.1 Delegation and Appropriate Reliance

Human-Computer Interaction (HCI) and Computer-Supported Cooperative Work (CSCW) research examines when and how people delegate decisions to AI systems and how reliance can be calibrated. Prior work proposes mechanisms such as *conditional delegation* [24] and studies whether explanations enable *complementary team performance* [6]. Related approaches formalize acceptance and override decisions [38] and introduce interventions to reduce overreliance, such as cognitive forcing functions [11]. More recent work highlights contextual drivers and cognitive barriers to effective delegation [17, 41]. While these frameworks advance task-level collaboration, they largely assume agency is preserved as long as users retain formal authority and override capability. **They do not explain why agency may nonetheless remain unexercised under sustained, convenience-driven delegation.**

CHI EA '26, Barcelona, Spain

ACM ISBN 979-8-4007-2281-3/26/04
https://doi.org/10.1145/3772363.3798760

2.2 Automation, Deskilling, and Skill Degradation

Human factors research documents automation-induced complacency, reduced monitoring, and skill degradation over time [34, 35]. Classic accounts of deskilling describe how automation substitutes human skills across domains [4, 5], with contemporary extensions to AI-assisted work and learning. However, this literature primarily conceptualizes agency erosion as *capability loss* through disuse. **It does not capture situations in which users retain capability but face a narrowed option space and increased friction when exercising choice.**

2.3 Infrastructure Studies and Path Dependency

Science and Technology Studies (STS) and HCI conceptualize infrastructure as relational and typically invisible in routine practice [42], emphasizing how standards and classifications encode choices that shape future action [10]. Path dependency theories explain lock-in through increasing returns and quasi-irreversible investments [2, 3, 13, 14]. Within HCI, infrastructure perspectives are increasingly applied to digital systems and user coping practices [26]. Yet these accounts focus primarily on macro-level systems, leaving the micro-level dynamics by which repeated interface-level delegations become infrastructural under-theorized.

2.4 Summary: Toward Decision Horizons

Together, these strands explain (i) task-level reliance decisions, (ii) capability degradation through disuse, and (iii) macro-level lock-in through path dependency. However, they do not account for a recurring pattern in LLM-mediated interaction: agency erosion as a gradual, cumulative outcome of repeated, rational delegation *even when formal control and capability remain intact*. Strand (i) assumes agency is preserved as long as users retain override authority; strand (ii) requires observable skill loss as a precondition; strand (iii) operates at systemic rather than interaction level. None captures how convenience-driven delegation can reshape decision-making conditions without coercion, capability loss, or explicit lock-in.

We address this gap by introducing *decision horizons*, defined as the set of options and framings that remain epistemically visible and cognitively available to a user at a given point in time. Decision horizons are distinct from available options in a formal sense: an option may exist without being salient, plausible, or accessible for deliberation. We further theorize LLMs as *delegated decision infrastructures* that, through repeated use, systematically reshape which options and framings remain visible over time. This process does not remove alternatives but renders them progressively less accessible. This distinguishes decision horizon narrowing from automation bias (a single-instance anchoring effect), deskilling (capability loss through disuse), and overreliance (miscalibrated trust in system outputs).

3 A Four-Layer Model of LLM-Mediated Delegation

This section introduces a four-layer model of LLM-mediated delegation, visualized in Figure 1. The layers are ordered by increasing impact on human agency, but are not strictly sequential. Rather, they represent analytically distinct layers of delegation that may coexist or shift over time depending on context and use.

3.1 Layer 1: Assisted Decision-Making

At Layer 1, LLMs function as assistive tools that generate options or recommendations without constraining decision-making. Users retain epistemic access to alternatives and can evaluate, compare, and override outputs at negligible cost.

This mode aligns with work on appropriate reliance [38] and complementary team performance [6]. While Layer 1 can be sustained through domain expertise and interface designs that avoid single defaults, it is unstable in practice. As trust increases and LLMs integrate into workflows, the perceived value of independent evaluation declines, creating drift toward deeper forms of delegation [18, 25].

3.2 Layer 2: Epistemic Framing

Layer 2 captures delegation through problem framing rather than explicit decision substitution. Rather than directly making decisions, LLMs shape what appears relevant, plausible, or worth considering, thereby influencing which options enter the user's decision horizon.

Decision horizons describe not merely available options, but their epistemic accessibility, the degree to which alternatives are salient, plausible, and cognitively available for deliberation. In conversational interaction, early responses anchor subsequent reasoning, making later alternatives less prominent. Framing effects emerge from training distributions and prompt formulations [39, 40], sequential refinement within an initially established frame [20], and the resolution of underspecified requests into singular interpretations [46].

Recent empirical work demonstrates that such framing effects are not merely theoretical. For example, Goh et al. [19] show in a randomized clinical trial that access to an LLM significantly influences physicians' diagnostic reasoning processes, even when it does not consistently improve decision quality.

Although users retain formal authority, reframing requires recognizing the dependency of the system's outputs and actively reconstructing the problem space. This metacognitive step is cognitively demanding and often bypassed [23]. Empirical work confirms this: Gerlich [18] documents a negative correlation between AI tool usage and critical thinking, mediated by cognitive offloading. Agency thus remains formally intact while decision horizons narrow through repeated exposure to system-generated framings, constraining conceivable alternatives.

3.3 Layer 3: Infrastructural Dependency

At Layer 3, LLMs become embedded as default interfaces within recurring workflows. Delegation shifts from episodic use to infrastructural reliance.

Infrastructure studies characterize such systems as relational and largely invisible in routine practice [10, 42]. Lock-in emerges through tool integration [7], network effects [15], and data accumulation that raises switching costs. Cognitive offloading further increases opt-out effort even before skills degrade [37]. Ferdman [16] terms such contexts "capacity-hostile environments," systemic

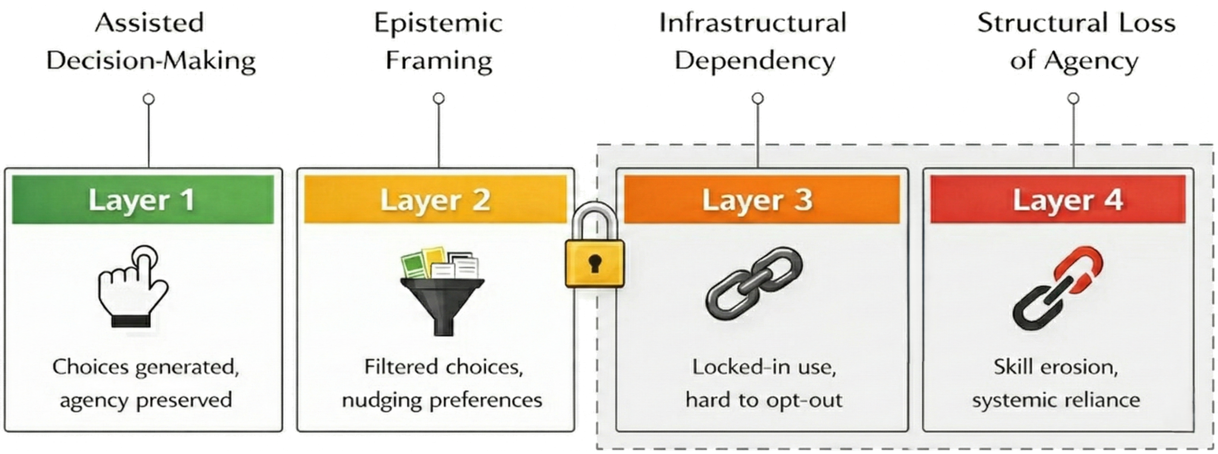

Figure 1: Four-layer model of LLM-mediated delegation. The model illustrates analytically distinct layers of delegation, ranging from assisted decision-making to structural loss of agency. Layers are ordered by increasing impact on agency but are not strictly sequential.

conditions that impede human capacity cultivation through AI mediation. Agency becomes latent: alternatives remain possible, but exercising them incurs growing friction.

3.4 Layer 4: Structural Loss of Agency

Layer 4 captures conditions under which agency becomes structurally eroded. Sustained delegation can lead to skill atrophy [35, 45], with recovery requiring substantial retraining or becoming effectively irreversible [32]. In parallel, non-LLM infrastructures may disappear as institutions, workflows, and norms adapt to automated systems [12]. Messeri and Crockett [30] warn that AI tools may create "illusions of understanding" in which users believe they understand more than they do, a condition that may mask capability erosion until reversal becomes impractical.

At this layer, reversibility is minimal. Individually rational acts of delegation can aggregate into collective capability loss that no single actor intends or controls [33].

Recent evidence supports this dynamic. The Anthropic Economic Index Report 2026 [1], based on 2M Claude conversations, shows that AI is disproportionately used for higher-education tasks (mean 14.4 vs. 13.2 years of schooling). This pattern suggests that sustained delegation concentrates cognitive labor in AI systems, contributing to gradual deskilling across occupations.

3.5 Transitions Between Layers

The layers are modes rather than fixed stages. Routine tasks tend toward infrastructural dependency, while ambiguous or high-stakes tasks may remain at earlier layers [25]. Transitions are asymmetric: while early layers can be preserved by design, movement toward Layers 3 and 4 introduces structural barriers to return. This asymmetry motivates interventions that prevent problematic transitions rather than assuming reversibility.

Delegation does not inherently erode agency. At Layer 1, offloading routine or well-defined tasks to LLMs can expand the range of options a user considers, reduce cognitive load, and free capacity for higher-order judgment. Agency erosion is not a necessary consequence of delegation but an emergent risk that arises when delegation becomes habitual, unreflective, or infrastructurally entrenched. The model is therefore not a critique of delegation per se, but a framework for distinguishing conditions under which delegation preserves or undermines meaningful human agency.

4 Implications for Design and Research

The four-layer model provides a vocabulary for analyzing agency in LLM-mediated systems and for deriving design and research implications. We summarize design strategies by layer, followed by directions for research and governance.

4.1 Design Strategies by Layer

Layer 1: Preserve epistemic access. To sustain assisted decision-making, interfaces should preserve access to alternatives and avoid making LLM outputs the default. Strategies include presenting multiple options in parallel to reduce anchoring [44], surfacing uncertainty and model limitations [47], and supporting lightweight comparison with non-LLM sources.

Layer 2: Make framings visible. Mitigating epistemic framing requires exposing otherwise implicit assumptions. Interventions include prompting users to articulate their framing, providing contrastive outputs that reveal how alternative framings affect results [31], and supporting periodic reflection. Recent work on Human-AI Deliberation [27] demonstrates that structured discussion mechanisms enabling dimension-level opinion elicitation can counteract framing effects and foster appropriate reliance. While such measures can preserve agency, they introduce friction that may affect adoption.

Layer 3: Enable reversibility. To limit infrastructural lock-in, systems should support partial delegation and graceful exit. This includes maintaining compatibility with non-LLM workflows, enabling hybrid interaction modes [21], and ensuring data portability and interoperability.

Layer 4: Preserve capability. Preventing structural agency loss requires maintaining human skills during extended delegation. Approaches include periodic manual practice, adaptive automation that responds to risk of skill decay [8, 28, 34], and professional standards that define baseline competencies independent of tool availability.

4.2 Research Directions

At Layer 3, opt-out friction can be assessed through two complementary measures: task completion time without LLM access after sustained use (behavioral), and self-reported delegation cost ("How difficult would it be to complete this task without AI support?", longitudinal, Likert-scaled). Divergence between perceived and actual capability serves as an additional indicator of infrastructural lock-in.

At Layer 4, skill atrophy measures from human factors research apply directly, including retention testing and error rate monitoring after delegation periods. Crucially, Layer 4 indicators should be combined with Layer 2 and 3 measures to distinguish capability loss from horizon narrowing and friction effects.

Across layers, we propose a short self-report instrument to operationalize layer transitions. Example items include: *"I formulate problems differently than before using AI"* (Layer 2), *"I rarely verify AI outputs anymore"* (Layer 3), and *"I would not know how to approach this task without AI"* (Layer 4). Longitudinal administration of such an instrument would allow researchers to trace delegation trajectories and evaluate interventions aimed at preserving agency.

Domain specificity remains essential: epistemic framing effects may dominate in creative and knowledge work, while structural capability loss is critical in safety-sensitive contexts such as medicine and law. Domain-specific validation of these measures is therefore a priority for future work.

5 Discussion

A central insight of the four-layer model is that agency can erode even when responsibility formally remains with the human decision-maker. Unlike traditional automation, where control and responsibility are explicitly transferred, LLM-mediated delegation preserves formal accountability while degrading the practical conditions for exercising agency. This produces a responsibility gap [29], in which users are accountable for outcomes shaped by epistemic framings and infrastructures they do not fully control, particularly in high-stakes domains such as healthcare, law, and hiring [9].

The model clarifies how LLM-mediated delegation differs from existing human–AI frameworks. Whereas appropriate reliance and trust calibration optimize individual acceptance decisions and automation taxonomies emphasize levels of autonomy [36], the four-layer model captures cumulative effects emerging through repeated delegation. Distinguishing infrastructural dependency (Layer 3) from structural capability loss (Layer 4) highlights an intermediate, partially reversible stage of agency erosion that is obscured when dependency and deskilling are conflated.

Calibrated trust is necessary but insufficient for preserving agency. Even rational reliance can narrow decision horizons through epistemic framing (Layer 2) or increase exit costs through infrastructural integration (Layer 3). Tang [43] describes this as "epistemic cost," harm to knowledge quality arising from AI-mediated production even absent ethical violations. Preserving agency therefore requires addressing structural features of interaction and infrastructure, not only improving system accuracy or trust calibration.

At Layer 4, agency loss becomes a collective concern. When rational delegation aggregates into capability erosion, individually beneficial choices may yield socially undesirable outcomes [33]. Potential responses include transparency requirements, educational standards, interoperability mandates, and professional ethics codes that safeguard human judgment [22]. The aim is not to restrict delegation but to ensure reversibility and preserve meaningful agency.

Addressing cumulative delegation effects also requires metacognitive awareness. Users who recognize their delegation patterns and epistemic consequences are better positioned to preserve agency. Educational interventions that foreground such reflection may complement technical design strategies, particularly in higher education contexts where LLM use is concentrated [1].

This work has limitations. The model is conceptual and requires empirical validation, especially through longitudinal studies. It foregrounds cognitive and epistemic dimensions of agency, while material and socioeconomic factors remain only partially addressed. Finally, layer boundaries are analytically useful but may blur in practice.

6 Conclusion

Nowadays, LLMs are reshaping the conditions under which human decisions are made – no questions asked! We argue that understanding this shift requires moving beyond task-level performance and binary notions of control toward an analysis of how delegation accumulates, decision horizons narrow, and agency becomes latent over time.

The proposed four-layer model distinguishes assisted decision-making, epistemic framing, infrastructural dependency, and structural loss of agency as qualitatively different layers of LLM-mediated delegation. Together, these layers explain how agency erosion can emerge gradually through rational, convenience-driven use rather than through coercion, misalignment, or immediate skill loss.

Three implications follow. First, agency erosion is cumulative and emergent, arising from repeated delegations that reshape workflows and expectations. Second, reversibility is asymmetric: as interaction shifts toward infrastructural dependency and capability loss, opt-out costs rise and independent performance becomes harder to restore. Third, responsibility and agency can decouple, leaving users accountable for decisions shaped by conditions they did not deliberately choose.

For HCI research, this perspective motivates evaluation criteria beyond accuracy and efficiency, including decision horizons, opt-out friction, and capability preservation. For design, it foregrounds reversibility, framing transparency, and sustained human skill. For governance, it highlights how individually rational delegation may aggregate into collective capability loss.

While LLMs offer substantial value as assistive systems, the challenge is realizing that value without undermining meaningful human agency. This requires deliberate design, empirical scrutiny, and attention to how delegation reshapes not only what systems can do, but what humans remain able to do.

Acknowledgments

This research is funded by the Federal Ministry of Research, Technology and Space of Germany (16SV9597, 03DPC0716C).

References

[1] Ruth Appel, Maxim Massenkoff, Peter McCrory, Miles McCain, Ryan Heller, Tyler Neylon, Alex Tamkin, et al. 2026. *The Anthropic Economic Index Report: Economic Primitives*. Technical Report. Anthropic. https://www.anthropic.com/research/anthropic-economic-index-january-2026-report Published January 15, 2026.

[2] W. Brian Arthur. 1989. Competing Technologies, Increasing Returns, and Lock-In by Historical Events. *The Economic Journal* 99, 394 (1989), 116–131.

[3] W. Brian Arthur. 1994. *Increasing Returns and Path Dependence in the Economy*. University of Michigan Press, Ann Arbor.

[4] David H. Autor. 2015. Why Are There Still So Many Jobs? The History and Future of Workplace Automation. *Journal of Economic Perspectives* 29, 3 (2015), 3–30. doi:10.1257/jep.29.3.3

[5] Lisanne Bainbridge. 1983. Ironies of Automation. *Automatica* 19, 6 (1983), 775–779. doi:10.1016/0005-1098(83)90046-8

[6] Gagan Bansal, Tongshuang Wu, Joyce Zhou, Raymond Fok, Besmira Nushi, Ece Kamar, Marco Tulio Ribeiro, and Daniel Weld. 2021. Does the Whole Exceed its Parts? The Effect of AI Explanations on Complementary Team Performance. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM, 1–16. doi:10.1145/3411764.3445717

[7] Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (Virtual Event, Canada) *(FAccT '21)*. Association for Computing Machinery, New York, NY, USA, 610–623. doi:10.1145/3442188.3445922

[8] Margherita Bernabei and Francesco Costantino. 2024. Adaptive automation: Status of research and future challenges. *Robotics and Computer-Integrated Manufacturing* 88 (2024), 102724. doi:10.1016/j.rcim.2024.102724

[9] Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, Erik Brynjolfsson, S. Buch, Dallas Card, Rodrigo Castellon, Niladri S. Chatterji, Annie S. Chen, Kathleen A. Creel, Jared Davis, Dora Demszky, Chris Donahue, Moussa Koulako Bala Doumbouya, Esin Durmus, Stefano Ermon, John Etchemendy, Kawin Ethayarajh, Li Fei-Fei, Chelsea Finn, Trevor Gale, Lauren Gillespie, Karan Goel, Noah D. Goodman, Shelby Grossman, Neel Guha, Tatsunori Hashimoto, Peter Henderson, John Hewitt, Daniel E. Ho, Jenny Hong, Kyle Hsu, Jing Huang, Thomas F. Icard, Saahil Jain, Dan Jurafsky, Pratyusha Kalluri, Siddharth Karamcheti, Geoff Keeling, Fereshte Khani, O. Khattab, Pang Wei Koh, Mark S. Krass, Ranjay Krishna, Rohith Kuditipudi, Ananya Kumar, Faisal Ladhak, Mina Lee, Tony Lee, Jure Leskovec, Isabelle Levent, Xiang Lisa Li, Xuechen Li, Tengyu Ma, Ali Malik, Christopher D. Manning, Suvir Mirchandani, Eric Mitchell, Zanele Munyikwa, Suraj Nair, Avanika Narayan, Deepak Narayanan, Benjamin Newman, Allen Nie, Juan Carlos Niebles, Hamed Nilforoshan, Julian Nyarko, Giray Ogut, Laurel J. Orr, Isabel Papadimitriou, Joon Sung Park, Chris Piech, Eva Portelance, Christopher Potts, Aditi Raghunathan, Robert Reich, Hongyu Ren, Frieda Rong, Yusuf H. Roohani, Camilo Ruiz, Jack Ryan, Christopher R'e, Dorsa Sadigh, Shiori Sagawa, Keshav Santhanam, Andy Shih, Krishna Parasuram Srinivasan, Alex Tamkin, Rohan Taori, Armin W. Thomas, Florian Tramèr, Rose E. Wang, William Wang, Bohan Wu, Jiajun Wu, Yuhuai Wu, Sang Michael Xie, Michihiro Yasunaga, Jiaxuan You, Matei A. Zaharia, Michael Zhang, Tianyi Zhang, Xikun Zhang, Yuhui Zhang, Lucia Zheng, Kaitlyn Zhou, and Percy Liang. 2021. On the Opportunities and Risks of Foundation Models. *ArXiv* abs/2108.07258 (2021). https://api.semanticscholar.org/CorpusID:237091588

[10] Geoffrey C. Bowker and Susan Leigh Star. 1999. *Sorting Things Out: Classification and Its Consequences*. MIT Press, Cambridge, MA.

[11] Zana Buçinça, Maja Barbara Malaya, and Krzysztof Z. Gajos. 2021. To Trust or to Think: Cognitive Forcing Functions Can Reduce Overreliance on AI in AI-assisted Decision-making. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW1 (2021), 1–21. doi:10.1145/3449287

[12] Nicholas Carr. 2014. *The Glass Cage: Automation and Us*. W. W. Norton & Company, New York.

[13] Paul A. David. 1985. Clio and the Economics of QWERTY. *American Economic Review* 75, 2 (1985), 332–337.

[14] Paul A. David. 2007. Path Dependence: A Foundational Concept for Historical Social Science. *Cliometrica* 1, 2 (2007), 91–114. doi:10.1007/s11698-006-0005-x

[15] Nicholas Economides. 1996. The Economics of Networks. *International Journal of Industrial Organization* 14, 6 (1996), 673–699. doi:10.1016/0167-7187(96)01015-6

[16] Avigail Ferdman. 2025. AI deskilling is a structural problem. *AI & SOCIETY* (05 Nov 2025). doi:10.1007/s00146-025-02686-z

[17] Andreas Fügener, Jörn Grahl, Alok Gupta, and Wolfgang Ketter. 2022. Cognitive Challenges in Human–Artificial Intelligence Collaboration: Investigating the Path Toward Productive Delegation. *Information Systems Research* 33, 2 (2022), 678–696. doi:10.1287/isre.2021.1079

[18] Michael Gerlich. 2025. AI Tools in Society: Impacts on Cognitive Offloading and the Future of Critical Thinking. *Societies* 15, 1 (2025). doi:10.3390/soc15010006

[19] Ethan Goh, Robert Gallo, Jason Hom, Eric Strong, Yingjie Weng, Hannah Kerman, Joséphine A. Cool, Zahir Kanjee, Andrew S. Parsons, Neera Ahuja, Eric Horvitz, Daniel Yang, Arnold Milstein, Andrew P. J. Olson, Adam Rodman, and Jonathan H. Chen. 2024. Large Language Model Influence on Diagnostic Reasoning: A Randomized Clinical Trial. *JAMA Network Open* 7, 10 (10 2024), e2440969–e2440969. doi:10.1001/jamanetworkopen.2024.40969

[20] Jess Hohenstein, Rene F. Kizilcec, Dominic DiFranzo, Zhila Aghajari, Hannah Mieczkowski, Karen Levy, Mor Naaman, Jeffrey Hancock, and Malte F. Jung. 2023. Artificial Intelligence in Communication Impacts Language and Social Relationships. *Scientific Reports* 13, 1 (2023), 5487. doi:10.1038/s41598-023-30938-9

[21] Eric Horvitz. 1999. Principles of Mixed-Initiative User Interfaces. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (1999), 159–166. doi:10.1145/302979.303030

[22] Anna Jobin, Marcello Ienca, and Effy Vayena. 2019. The Global Landscape of AI Ethics Guidelines. *Nature Machine Intelligence* 1, 9 (2019), 389–399. doi:10.1038/s42256-019-0088-2

[23] Daniel Kahneman. 2011. *Thinking, Fast and Slow*. Farrar, Straus and Giroux, New York.

[24] Vivian Lai, Samuel Carton, Rajat Bhatnagar, Q. Vera Liao, Yunfeng Zhang, and Chenhao Tan. 2022. Human-AI Collaboration via Conditional Delegation: A Case Study of Content Moderation. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. ACM, 1–18. doi:10.1145/3491102.3501999

[25] John D Lee and Katrina A See. 2004. Trust in automation: designing for appropriate reliance. *Human factors* 46, 1 (2004), 50–80. doi:10.1518/hfes.46.1.50_30392

[26] Yao Lyu, Jie Cai, and John M. Carroll. 2025. A Systematic Literature Review of Infrastructure Studies in SIGCHI. *Proc. ACM Hum.-Comput. Interact.* 9, 7, Article CSCW364 (Oct. 2025), 41 pages. doi:10.1145/3757545

[27] Shuai Ma, Qiaoyi Chen, Xinru Wang, Chengbo Zheng, Zhenhui Peng, Ming Yin, and Xiaojuan Ma. 2025. Towards Human-AI Deliberation: Design and Evaluation of LLM-Empowered Deliberative AI for AI-Assisted Decision-Making. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 261, 23 pages. doi:10.1145/3706598.3713423

[28] Brooke N Macnamara, Ibrahim Berber, M Cenk Çavuşoğlu, Elizabeth A Krupinski, Naren Nallapareddy, Noelle E Nelson, Philip J Smith, Amy L Wilson-Delfosse, and Soumya Ray. 2024. Does using artificial intelligence assistance accelerate skill decay and hinder skill development without performers' awareness? *Cogn. Res. Princ. Implic.* 9, 1 (July 2024), 46.

[29] Andreas Matthias. 2004. The Responsibility Gap: Ascribing Responsibility for the Actions of Learning Automata. *Ethics and Information Technology* 6, 3 (2004), 175–183. doi:10.1007/s10676-004-3422-1

[30] Lisa Messeri and M. J. Crockett. 2024. Artificial intelligence and illusions of understanding in scientific research. *Nature* 627, 8002 (01 Mar 2024), 49–58. doi:10.1038/s41586-024-07146-0

[31] Tim Miller. 2019. Explanation in Artificial Intelligence: Insights from the Social Sciences. *Artificial Intelligence* 267 (2019), 1–38. doi:10.1016/j.artint.2018.07.007

[32] Chiara Natali, Luca Marconi, Leslye Denisse Dias Duran, and Federico Cabitza. 2025. AI-induced Deskilling in Medicine: A Mixed-Method Review and Research Agenda for Healthcare and Beyond. *Artificial Intelligence Review* 58, 11 (27 Aug 2025), 356. doi:10.1007/s10462-025-11352-1

[33] Elinor Ostrom. 1990. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press.

[34] Raja Parasuraman and Dietrich H. Manzey. 2010. Complacency and Bias in Human Use of Automation: An Attentional Integration. *Human Factors* 52, 3 (2010), 381–410. doi:10.1177/0018720810376055

[35] Raja Parasuraman and Victor Riley. 1997. Humans and Automation: Use, Misuse, Disuse, Abuse. *Human Factors* 39, 2 (1997), 230–253. doi:10.1518/001872097778543886

[36] Raja Parasuraman, Thomas B. Sheridan, and Christopher D. Wickens. 2000. A Model for Types and Levels of Human Interaction with Automation. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans* 30, 3 (2000), 286–297. doi:10.1109/3468.844354

[37] Evan F. Risko and Sam J. Gilbert. 2016. Cognitive Offloading. *Trends in Cognitive Sciences* 20, 9 (2016), 676–688. doi:10.1016/j.tics.2016.07.002

[38] Max Schemmer, Niklas Kuehl, Carina Benz, Andrea Bartos, and Gerhard Satzger. 2023. Appropriate Reliance on AI Advice: Conceptualization and the Effect of

Explanations. In *Proceedings of the 28th International Conference on Intelligent User Interfaces*. ACM, 410–422. doi:10.1145/3581641.3584066

[39] Melanie Sclar, Yejin Choi, Yulia Tsvetkov, and Alane Suhr. 2024. Quantifying Language Models' Sensitivity to Spurious Features in Prompt Design or: How I learned to start worrying about prompt formatting. In *The Twelfth International Conference on Learning Representations*. https://openreview.net/forum?id=RIu5lyNXjT

[40] Murray Shanahan, Kyle McDonell, and Laria Reynolds. 2023. Role Play with Large Language Models. *Nature* 623 (2023), 493–498. doi:10.1038/s41586-023-06647-8

[41] Philipp Spitzer, Joshua Holstein, Patrick Hemmer, Michael Vössing, Niklas Kühl, Dominik Martin, and Gerhard Satzger. 2025. Human Delegation Behavior in Human-AI Collaboration: The Effect of Contextual Information. *Proceedings of the ACM on Human-Computer Interaction* 8, CSCW (2025), 1–39. doi:10.1145/3710999

[42] Susan Leigh Star and Karen Ruhleder. 1996. Steps Toward an Ecology of Infrastructure: Design and Access for Large Information Spaces. In *Information Systems Research*, Vol. 7. INFORMS, 111–134. doi:10.1287/isre.7.1.111

[43] Bor Luen Tang. 2025. The Epistemic Downside of Using LLM-Based Generative AI in Academic Writing. *Publications* 13, 4 (2025). doi:10.3390/publications13040063

[44] Amos Tversky and Daniel Kahneman. 1974. Judgment under Uncertainty: Heuristics and Biases. *Science* 185, 4157 (1974), 1124–1131. doi:10.1126/science.185.4157.1124

[45] Lei Wang, Chen Ma, Xueyang Feng, Zeyu Zhang, Hao Yang, Jingsen Zhang, Zhiyuan Chen, Jiakai Tang, Xu Chen, Yankai Lin, Wayne Xin Zhao, Zhewei Wei, and Jirong Wen. 2024. A survey on large language model based autonomous agents. *Frontiers of Computer Science* 18, 6 (22 Mar 2024), 186345. doi:10.1007/s11704-024-40231-1

[46] Laura Weidinger, Jonathan Uesato, Maribeth Rauh, Conor Griffin, Po-Sen Huang, John Mellor, Amelia Glaese, Myra Cheng, Borja Balle, Atoosa Kasirzadeh, et al. 2022. Taxonomy of Risks Posed by Language Models. *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency* (2022), 214–229. doi:10.1145/3531146.3533088

[47] Yunfeng Zhang, Q. Vera Liao, and Rachel K. E. Bellamy. 2020. Effect of Confidence and Explanation on Accuracy and Trust Calibration in AI-Assisted Decision Making. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (2020), 295–305. doi:10.1145/3351095.3372852