

Broader Impacts of GenAI in Communication: Building Agendas for Research, Design, and Policymaking

Jiayin Zhi
jzhi@uchicago.edu
University of Chicago
Chicago, IL, USA

Lan Gao*
langao@uchicago.edu
University of Chicago
Chicago, IL, USA

Angel Hsing-Chi Hwang
angel.hwang@usc.edu
University of South California
Los Angeles, CA, USA

Jessica He
jessicahe@ibm.com
IBM Research
Cambridge, MA, USA

Alexis Hiniker
alexisr@uw.edu
University of Washington
Seattle, WA, USA

Munmun De Choudhury
munmund@gatech.edu
Georgia Institute of Technology
Atlanta, GA, USA

Mina Lee
mnlee@cs.uchicago.edu
University of Chicago
Chicago, IL, USA

Marshini Chetty
marshini@uchicago.edu
University of Chicago
Chicago, IL, USA

Abstract

GenAI is increasingly integrated into a broad range of activities in communication, from drafting messages and creating content to summarizing and filtering information that people receive. It has also become commonly used across a wide range of settings, from interpersonal to social. Accordingly, the impacts of GenAI have extended well beyond each communication session, which prior research has primarily focused on, and warrant a more holistic examination. For example: How can GenAI's mediation of language affect human cognition and well-being? How can the integration of GenAI in communication shift people's opinions and beliefs? How should ownership and attribution of AI-generated content be defined? How can the erosion of effort and authenticity signals in communication by GenAI affect trust and relationships? Many of these questions remain largely open. **More importantly, we lack a shared agenda for whether and how these impacts constitute risks that warrant action, and what research, design, and policy responses are needed.** We call on the CSCW community—bringing together researchers, designers, and practitioners—to build empirical foundations for understanding the range of these impacts, identify a broader set of corresponding risks, and translate findings into concrete design and policy agendas for mitigation.

CCS Concepts

• **Human-centered computing** → **Collaborative and social computing systems and tools.**

Keywords

AI-mediated communication, generative AI, AI risks, AI safety

*Corresponding author.



This work is licensed under a Creative Commons Attribution 4.0 International License.
CSCW Companion '26, Salt Lake City, UT, USA
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2378-0/2026/10
<https://doi.org/10.1145/3785651.3816118>

ACM Reference Format:

Jiayin Zhi, Lan Gao, Angel Hsing-Chi Hwang, Jessica He, Alexis Hiniker, Munmun De Choudhury, Mina Lee, and Marshini Chetty. 2026. Broader Impacts of GenAI in Communication: Building Agendas for Research, Design, and Policymaking. In *Companion of the Computer-Supported Cooperative Work and Social Computing (CSCW Companion '26)*, October 10–14, 2026, Salt Lake City, UT, USA. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3785651.3816118>

1 Introduction

As generative AI (GenAI) has become an everyday application, AI-mediated communication—“mediated communication between people in which a computational agent operates on behalf of a communicator by modifying, augmenting, or generating messages to accomplish communication or interpersonal goals” [6]—has become an increasingly common practice [7]. GenAI has now been integrated into a broad range of communication activities in various interpersonal and social settings. On the side of communicators, to name a few examples, people now use GenAI to generate messages such as emails and social network comments [4, 6], co-write content for professional communication and public discourse [16, 20], and create images, audio, and videos to share ideas and express themselves online [16, 20]. On the side of receivers, GenAI increasingly mediates how people consume communication—for instance, summarizing large volumes of text [1, 31], assisting reading and comprehension [34], and filtering what information people encounter in their feeds and search results [22, 27].

The communication activities that GenAI is integrated into encompass comprehensive processes of writing, creating, reading, understanding, and reasoning, through which people express, exchange, and make sense of information and ideas [9, 29]. Such integrations involve human capacities that matter not only for sustaining individual development [28, 33], but also for maintaining community and social bonds [10, 14, 23, 24], collective knowledge [30], and public discourse [2]. Prior research examining AI's impact on AI-mediated communication has primarily focused on the experiences and outcomes of a single communication session, such

as examining receiver trust over AI-generated communicative materials [17], how AI-generated messages persuade or even mislead the receiver [18, 35], and perceptions of the message sender on how the AI agent represents them [14]. However, given the increasing integration of GenAI into a wider range of activities in communication across interpersonal and social settings, **the range of impacts broadens correspondingly**.

Specifically, these impacts span several interconnected dimensions. At the individual level, as communication activities are the means through which people exercise and develop cognition, questions arise about how GenAI's mediation influences individuals' cognitive capacities, such as creativity in content creation and critical thinking in consuming communication [19, 28, 33]. At the interpersonal level, research has shown that AI-mediated communication may erode signals of authenticity and effort that sustain trust and relationships in interpersonal communication and online communities [10, 13, 14, 23, 24]. For public discourse, there is growing evidence that GenAI integration can subtly shape people's attitudes, and may eventually change their beliefs in ways that may contribute to the polarization of opinions [2, 16]. Furthermore, as AI-generated content becomes an increasingly prevalent form of communicative material, it introduces legal and ethical debates: questions of ownership and attribution—who authored a message, post, or piece of content—affect the reception of communication, yet no consensus has been reached for now [3, 8, 15, 21]. And the proliferation of AI slop—low-effort, mass-produced AI-generated content flooding communication channels—raises open questions about human judgment and taste of communicative content, and the recognition and valuation of human creative work [21, 26]. In most cases, the risks arising from such impacts could be hard to attribute to a single bad actor or deployment decision, and may only become apparent as they accumulate gradually across many subtle instances. For instance, people's epistemic autonomy may be eroded if they increasingly rely on GenAI for searching and summarizing communicative materials [25]; and the trust in online communities may be undermined due to the increasing loss of authenticity in communication as AI-generated content becomes indistinguishable [5, 23, 24].

Researchers have begun to investigate these new impacts and develop corresponding design solutions to scaffold GenAI use, yet many of these questions remain largely open. More importantly, **we lack a shared agenda for whether and how these impacts constitute risks that warrant action, and what research, design, and policy responses are needed**. To broaden the conversation on these open challenges and advance cross-community agendas for the future of AI-mediated communication, this workshop calls on the CSCW community—bringing together researchers, designers, and practitioners—to share empirical evidence for understanding the range of these impacts, surface questions that require more attention or are less investigated, develop methodologies to better examine such questions, identify a broader set of corresponding risks, brainstorm potential design solutions to mitigate or prevent those risks, and discuss what policy and regulatory responses may be needed.

2 Workshop Themes

This workshop aims to draw attention to and build agendas of research, design, and policies for emerging risks of GenAI use in communication across interpersonal and social settings—a gray area that remains underexamined and underaddressed. We welcome participants who are interested in or actively working on these open challenges, including those conducting empirical investigations of GenAI use in communication, those designing GenAI-integrated tools and systems that mediate communication, and those engaged in policy and governance work related to GenAI and communication.

To broaden the conversations and foster cross-community discussions, through structured working sessions in our workshop on the following themes, participants will share their thoughts or findings on the impacts of GenAI use in communication (**empirical foundations**), brainstorm and discuss design interventions that can scaffold GenAI in communication (**design opportunities and implications**), and discuss and identify aspects requiring policy and regulatory intervention (**policy and regulation**), particularly those that cannot be addressed through technical solutions and design alone. By convening across disciplines and sectors, we aim to build a shared vocabulary and chart concrete directions for research, design, and policymaking. Open questions we will cover in our workshop include but not limited to:

- **Empirical foundations:** What do we know about the impacts of GenAI in communication? What are the potential risks arising from these impacts? How do such impacts and risks differ across different settings, such as messaging within close relationships and professional communication? What research methods and frameworks are effective for studying these impacts and emerging risks?
- **Design opportunities and implications:** How can GenAI tools (e.g., writing assistants and conversational agents) be designed to prevent the potential risks? How can digital platforms (e.g., messaging apps and social media platforms) be designed to respond to the rise of AI-mediated communication as GenAI adoption increases? What designs and mechanisms can help preserve meaningful human engagement while mitigating risks in communication with GenAI?
- **Policy and regulation:** What new policies and regulations are needed? What can or should be considered under the governance frameworks? How should debatable topics, such as ownership and attribution of AI-generated content, be regulated in communication? And how do we create accountability structures for corresponding consequences that are hard to attribute to any single stakeholder?

3 Workshop Activities

This one-day in-person workshop is structured around three progressive working sessions that move from identifying impacts and emergent risks, to exploring design responses, to discussing policy and regulatory implications. To foster early community building, participants will be invited to a Discord server two weeks before the workshop day. They will be asked to introduce themselves in a general channel. This Discord server serves as a starting point

Table 1: Tentative schedule of workshop activities, including asynchronous activities prior to the workshop day.

Duration	Session
<i>Pre-Workshop</i>	
2 weeks before	Participants introduce themselves in the Discord channel
<i>On the Workshop Day</i>	
30 min	Opening Remarks and Keynote Speech
50 min	Lightning Talks
15 min	<i>Coffee Break</i>
50 min	Working Session 1: Empirical Foundations
15 min	<i>Coffee Break</i>
50 min	Working Session 2: Design Opportunities and Implications
60 min	<i>Lunch Break</i>
50 min	Working Session 3: Policy and Regulation
15 min	<i>Coffee Break</i>
30 min	Plenary Reflection and Follow-Up Planning
15 min	Closing and Final Remarks
<i>Post-Workshop Social Activities (Optional)</i>	

for discussion ahead of the workshop and a space for follow-up collaboration afterward.

- **Opening Remarks and Keynote Speech (30 minutes):** Organizers introduce the workshop's goals, structure, and themes. A keynote speaker from our organizing team will offer an opening provocation to ground the day's discussions in the current landscape of GenAI in communication.
- **Lightning Talks (50 minutes):** Each participant will give a one-minute lightning talk introducing themselves and sharing one problem, idea, or insight from their work related to the workshop themes. Talks will be accompanied by one to two slides loaded into a single auto-advancing deck. This session is designed to help participants get to know one another and seed the discussions in the co-working sessions that follow.
- **Working Session 1: Empirical Foundations (50 minutes):** This session aims to collectively map what we know about the impacts of GenAI in communication, what risks may arise from these impacts, and what remains open and understudied. Prior to the workshop, organizers will synthesize key impacts and open questions from participants' submissions, dedicating one slide to each. During the session, participants will discuss each slide together—for instance, whether a given impact constitutes a risk, what evidence exists, how it varies across settings, and what methods are needed to investigate it further. This discussion will cover the open questions outlined in Section 2, collectively categorize the impacts based on stakes and types, and build a shared map of the impact landscape that will guide the remaining sessions.
- **Working Session 2: Design Opportunities and Implications (50 minutes):** Based on the impact landscape and higher-level categories identified in previous session, participants will self-organize into small groups aligned with

the categories based on their own interests. Each group will brainstorm and discuss design solutions that could scaffold GenAI use in communication and mitigate the identified risks within their category. In the final 15 minutes, each group will briefly share its ideas and reflect on what design can and cannot address, surfacing remaining concerns that may require responses beyond design.

- **Working Session 3: Policy and Regulation (50 minutes):** This session opens up the conversation to policy and regulatory responses. Drawing on the concerns and design limitations surfaced in the previous sessions, participants will discuss what responsibilities platforms, AI developers, and policymakers have, what can or should be regulated, and what governance mechanisms are needed to address harms that are gradual, distributed, and hard to attribute to any single actor.
- **Plenary Reflection and Follow-Up Planning (30 minutes):** Participants reconvene to reflect on key insights across all three sessions, identify priorities for future research, design, and policy agendas, and discuss opportunities for cross-community collaborations and follow-up works.
- **Closing and Final Remarks (15 minutes):** Organizers recap the day's key takeaways, invite participants interested in continued collaboration to connect, and outline next steps for any public dissemination of workshop outcomes. All contributions will be formally acknowledged.

4 Workshop Logistics

4.1 Pre-Workshop

For early community building, we will create a Discord server two weeks before the workshop and actively moderate the server to prompt discussion and offer support. We will also host a workshop

website with the call for participation, dates, agenda, logistics, and submission guidelines.

4.2 On the Workshop Day

We propose a one-day in-person workshop. For the lightning talks, participants will be asked to submit one to two slides in advance, which will be loaded into a single auto-advancing deck for the presentation. The workshop organizers will bring laptops for the presentation, as well as physical supplies including sticky notes, paper, cardboard, pens, and markers as alternative means for affinity diagramming, note-taking, brainstorming, and group discussions. Participants are expected to bring their own laptops to engage with online collaborative tools such as Miro Board and Google Doc for real-time notetaking and brainstorming during working sessions.

5 Call for Participants

We invite 30–50 participants who are engaged with or interested in topics of impacts and emerging risks of GenAI use in communication. We welcome participants from various academic and professional backgrounds who are interested in or actively working on these open challenges, including those conducting empirical investigations of GenAI use in communication, those designing GenAI-integrated tools and systems that mediate communication, and those engaged in policy and governance work related to GenAI and communication. Participants can apply through one of two submission formats:

- **Short Paper:** Participants are invited to submit a short paper (up to 4 pages, 2-column ACM format, excluding references) presenting their completed or ongoing works that are related to the impacts of GenAI use in communication. Submissions can reflect work in progress. All submissions will be reviewed by the organizing team, with selection guided by our commitment to assembling a group that reflects a diversity of disciplines, experiences, and perspectives. Accepted short papers will be published on our workshop website.
- **Statement of Interest:** Interested individuals who do not have a related paper to submit are invited to submit a brief statement (1–2 pages) expressing their interests in the impacts of GenAI use in communication with concrete examples, outlining their reasons for wanting to join the workshop, questions they want to explore, and a short bio highlighting their background or current work.

6 Workshop Outcomes

We aim to foster a cross-disciplinary community of researchers, designers, and practitioners. Through the working sessions and facilitated discussions, we hope to spark conversations and lay the groundwork for future collaborations and research partnerships. We will designate note takers for all working sessions to ensure detailed documentation of insights and recommendations. After the workshop, the organizing team will synthesize these notes into a summary that highlights key impact and risks, design opportunities, policy implications, and future research directions. The synthesized workshop summary will be shared on the workshop website for broader visibility, and we will consider submitting follow-up works

of the outcomes to a conference or journal to further advance community discussion. During the follow-up planning session of the workshop, we will invite those interested in continued collaboration to contribute to the follow-up works.

7 Organizing Team

- **Jiayin Zhi** is a PhD student in Computer Science at the University of Chicago. Her research investigates how GenAI influences the way people communicate and think by creating and evaluating GenAI systems. She co-organized a Special Interest Group at CSCW 2025 [32].
- **Lan Gao** is a PhD student in Computer Science at the University of Chicago. Her research focuses on the governance of GenAI in the online information ecosystem, such as AI-generated content in online communities and digital platforms.
- **Angel Hsing-Chi Hwang** is an Assistant Professor of Communication and Computer Science at the University of Southern California. Her research examines how conversational AIs mediate social settings and the impact they have. She has led and organized several workshops and panels at CSCW, CHI, and CHIWORK [3, 11, 12].
- **Jessica He** is a UX Designer at IBM Research, where she works on understanding user needs and risks in emerging AI technology to design safer, more transparent systems. Her work has covered topics such as ownership and attribution in co-created work and the integration of GenAI into collaborative workflows.
- **Alexis Hiniker** is an Associate Professor at the University of Washington Information School where she studies manipulative design and co-directs the Center for Digital Youth. She has provided subject-matter expertise to the White House, the U.S. Department of Justice, and the European Commission. She is a former scientific advisor to the U.S. Federal Trade Commission and a recipient of the SIGCHI Societal Impact award.
- **Munmun De Choudhury** is J. Z. Liang Professor at Georgia Tech, whose work sits at the intersection of computational social science, HCI, and digital mental health. She develops data-driven approaches to understand and support mental well-being at scale, and contributes to global AI and public health policy through advisory roles, including with the World Health Organization and other national and global initiatives. She is a recipient of the SIGCHI Societal Impact Award, a member of the SIGCHI Academy, and an ACM Distinguished Member.
- **Mina Lee** is an Assistant Professor in Computer Science and Data Science Institute at the University of Chicago. Her research examines writing, reading, and thinking with GenAI. She has co-founded and organized multiple workshops, including Intelligent and Interactive Writing Assistants (In2Writing), Human-centered Evaluation and Auditing of Language Models (HEAL), and Tools for Thought (TfT) at ACL and CHI.
- **Marshini Chetty** is an Associate Professor in the Department of Computer Science at the University of Chicago

where she directs the Amyoli Internet Research Laboratory (AIR lab). She specializes in human-computer interaction and usable privacy and security and has been actively studying how AI is being adopted in schools, scientific organizations, and AI governance more broadly.

References

- [1] Elliott Ash, Aniket Kesari, Suresh Naidu, Lena Song, and Dominik Stammbach. 2024. Translating legalese: enhancing public understanding of court opinions with legal summarizers. In *Proceedings of the 2024 Symposium on Computer Science and Law*. 136–157. doi:10.1145/3614407.3643700
- [2] Isabel Corpus, Eric Gilbert, Allison Koenecke, and Mor Naaman. 2025. Introducing AI to an Online Petition Platform Changed Outputs but not Outcomes. *arXiv preprint arXiv:2511.13949* (2025). doi:10.48550/arXiv.2511.13949
- [3] Hyo Jin Do, Molly Q Feldman, Jessica He, Angel Hsing-Chi Hwang, and Seyun Kim. 2025. Navigating Generative AI Disclosure, Ownership, and Accountability in Co-Creative Domains. In *Adjunct Proceedings of the 4th Annual Symposium on Human-Computer Interaction for Work (CHIWORK '25 Adjunct)*. Association for Computing Machinery, New York, NY, USA, Article 2, 4 pages. doi:10.1145/3707640.3729212
- [4] Yue Fu, Yixin Chen, Zelia Gomes Da Costa Lai, and Alexis Hiniker. 2025. Should ChatGPT write your breakup text? Exploring the role of AI in relationship dissolution. *Proceedings of the ACM on Human-Computer Interaction* 9, 7 (2025), 1–31. doi:10.1145/3757424
- [5] Lan Gao, Abani Ahmed, Oscar Chen, Margaux Reyl, Zayna Cheema, Nick Feamster, Chenhao Tan, Kurt Thomas, and Marshini Chetty. 2026. Governance of AI-Generated Content: A Case Study on Social Media Platforms. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*. 1–22. doi:10.1145/3772318.3790415
- [6] Jeffrey T Hancock, Mor Naaman, and Karen Levy. 2020. AI-mediated communication: Definition, research agenda, and ethical considerations. *Journal of Computer-Mediated Communication* 25, 1 (2020), 89–100. doi:10.1093/jcmc/zmz022
- [7] Kunal Handa, Alex Tamkin, Miles McCain, Saffron Huang, Esin Durmus, Sarah Heck, Jared Mueller, Jerry Hong, Stuart Ritchie, Tim Belonax, et al. 2025. Which economic tasks are performed with ai? evidence from millions of claude conversations. *arXiv preprint arXiv:2503.04761* (2025). doi:10.48550/arXiv.2503.04761
- [8] Jessica He, Stephanie Houde, and Justin D. Weisz. 2025. Which Contributions Deserve Credit? Perceptions of Attribution in Human-AI Co-Creation (CHI '25). Association for Computing Machinery, New York, NY, USA, Article 540, 18 pages. doi:10.1145/3706598.3713522
- [9] Susan C Herring. 2002. Computer-mediated. *Annual review of information science and technology* 109 (2002).
- [10] Jess Hohenstein and Malte Jung. 2020. AI as a moral crumple zone: The effects of AI-mediated communication on attribution and trust. *Computers in Human Behavior* 106 (2020), 106190. doi:10.1016/j.chb.2019.106190
- [11] Angel Hsing-Chi Hwang, Marianne Aubin Le Quéré, Hope Schroeder, Alejandro Cuevas, Steven P. Dow, Shivani Kapania, and Eugenia Rho. 2025. Scale, Engage, or Both?: Potential and Perils of Applying Large Language Models in Interview and Conversation-Based Research. In *Companion Publication of the 2025 Conference on Computer-Supported Cooperative Work and Social Computing (CSCW Companion '25)*. Association for Computing Machinery, New York, NY, USA, 95–98. doi:10.1145/3715070.3748284
- [12] Angel Hsing-Chi Hwang, Michael S. Bernstein, S. Shyam Sundar, Renwen Zhang, Manoel Horta Ribeiro, Yingdan Lu, Serina Chang, Tongshuang Wu, Aimei Yang, Dmitri Williams, Joon Sung Park, Katherine Ognyanova, Ziang Xiao, Aaron Shaw, and David A. Shamma. 2025. Human Subjects Research in the Age of Generative AI: Opportunities and Challenges of Applying LLM-Simulated Data to HCI Studies. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '25)*. Association for Computing Machinery, New York, NY, USA, Article 761, 7 pages. doi:10.1145/3706599.3716299
- [13] Angel Hsing-Chi Hwang, Q Vera Liao, Su Lin Blodgett, Alexandra Olteanu, and Adam Trischler. 2025. 'It was 80% me, 20% AI': Seeking Authenticity in Co-Writing with Large Language Models. *Proceedings of the ACM on Human-Computer Interaction* 9, 2 (2025), 1–41. doi:10.1145/3711020
- [14] Angel Hsing-Chi Hwang, John Oliver Siy, Renee Shelby, and Alison Lentz. 2024. In whose voice?: examining AI agent representation of people in social interaction through generative speech. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference*. 224–245. doi:10.1145/3643834.3661555
- [15] Angel Hsing-Chi Hwang, Senya Wong, Baixiao Chen, Jessica He, and Hyo Jin Do. 2026. "Better Ask for Forgiveness than Permission": Practices and Policies of AI Disclosure in Freelance Work. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*. 1–16. doi:10.1145/3772318.3791920
- [16] Maurice Jakesch, Advait Bhat, Daniel Buschek, Lior Zalmanson, and Mor Naaman. 2023. Co-writing with opinionated language models affects users' views. In *Proceedings of the 2023 CHI conference on human factors in computing systems*. 1–15. doi:10.1145/3544548.3581196
- [17] Maurice Jakesch, Megan French, Xiao Ma, Jeffrey T Hancock, and Mor Naaman. 2019. AI-mediated communication: How the perception that profile text was written by AI affects trustworthiness. In *Proceedings of the 2019 CHI conference on human factors in computing systems*. 1–13. doi:10.1145/3290605.3300469
- [18] Elise Karinshak, Sunny Xun Liu, Joon Sung Park, and Jeffrey T Hancock. 2023. Working with AI to persuade: Examining a large language model's ability to generate pro-vaccination messages. *Proceedings of the ACM on Human-Computer Interaction* 7, CSCW1 (2023), 1–29. doi:10.1145/3579592
- [19] Harsh Kumar, Jonathan Vincentius, Ewan Jordan, and Ashton Anderson. 2025. Human creativity in the age of llms: Randomized experiments on divergent and convergent thinking. In *Proceedings of the 2025 CHI conference on human factors in computing systems*. 1–18. doi:10.1145/3706598.3714198
- [20] Mina Lee, Katy Ilonka Gero, John Joon Yung Chung, Simon Buckingham Shum, Vipul Raheja, Hua Shen, Subhashini Venugopalan, Thiemo Wambagsans, David Zhou, Emad A Alghamdi, et al. 2024. A design space for intelligent and interactive writing assistants. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–35. doi:10.1145/3613904.3642697
- [21] Gabriel Lima, Nina Grgić-Hlača, and Elissa M. Redmiles. 2025. Public Opinions About Copyright for AI-Generated Art: The Role of Egocentricity, Competition, and Experience. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 1120, 32 pages. doi:10.1145/3706598.3713338
- [22] Anna Yan Liu and Maggie Mengqing Zhang. 2026. Positivity Bias in AI-Generated Summaries of User-Generated Content: Exploring Its Sources and Impact on Public Sentiment. *Computational Communication Research* 8, 2 (2026), 1. doi:10.5117/CCR2026.2.4.LIU
- [23] Travis Lloyd, Jennah Gosciak, Tung Nguyen, and Mor Naaman. 2025. AI rules? characterizing reddit community policies towards AI-generated content. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–19. doi:10.1145/3706598.3713292
- [24] Travis Lloyd, Joseph Reagle, and Mor Naaman. 2025. 'There Has To Be a Lot That We're Missing': Moderating AI-Generated Content on Reddit. *Proceedings of the ACM on Human-Computer Interaction* 9, 7 (2025), 1–24. doi:10.1145/3757445
- [25] Barbara Oakley, Michael Johnston, Ken-Zen Chen, Eulho Jung, and Terrence J Sejnowski. 2025. The memory paradox: Why our brains need knowledge in an age of AI. *arXiv preprint arXiv:2506.11015* (2025). doi:10.48550/arXiv.2506.11015
- [26] Irene Rae. 2024. The Effects of Perceived AI Use on Content Perceptions. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (Honolulu, HI, USA) (CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 978, 14 pages. doi:10.1145/3613904.3642076
- [27] 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*. 1–17. doi:10.1145/3613904.3642459
- [28] Anjali Singh, Karan Taneja, Zhitong Guan, and Avijit Ghosh. 2025. Protecting Human Cognition in the Age of AI. *arXiv:2502.12447 [cs.CY]* doi:10.48550/arXiv.2502.12447
- [29] Crispin Thurlow, Alice Tomic, and Lara Lengel. 2004. Computer mediated communication. (2004).
- [30] Sterling Williams-Ceci, Maurice Jakesch, Advait Bhat, Kowe Kadoma, Lior Zalmanson, and Mor Naaman. 2026. Biased AI writing assistants shift users' attitudes on societal issues. *Science Advances* 12, 11 (2026), eadw5578. doi:10.1126/sciadv.adw5578
- [31] Ziang Xiao, Tiffany Wenting Li, Karrie Karahalios, and Hari Sundaram. 2023. Inform the uninformed: improving online informed consent reading with an AI-powered Chatbot. In *Proceedings of the 2023 CHI Conference on human factors in computing systems*. 1–17. doi:10.1145/3544548.3581252
- [32] Alice Qian Zhang, Jiayin Zhi, Sravya Chandhiramowuli, Hong Shen, Laura Dabbish, Theodora Skeadas, Sarah Amos, and Jina Suh. 2025. The Work of AI Red Teaming: Automation and the Human Infrastructure. In *Companion Publication of the 2025 Conference on Computer-Supported Cooperative Work and Social Computing*. 84–87. doi:10.1145/3715070.3748281
- [33] Jiayin Zhi, Harsh Kumar, and Mina Lee. 2026. Investigating the Effects of LLM Use on Critical Thinking Under Time Constraints: Access Timing and Time Availability. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (Barcelona, Spain) (CHI '26)*. ACM, New York, NY, USA, 1–21. doi:10.1145/3772318.3791796
- [34] Jiayin Zhi, Hoyt Long, Richard Jean So, and Mina Lee. 2026. What Does AI Do for Cultural Interpretation? A Randomized Experiment on Close Reading Poems with Exposure to AI Interpretation. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (Barcelona, Spain) (CHI '26)*. ACM, New York, NY, USA, 1–18. doi:10.1145/3772318.3791727
- [35] Jiawei Zhou, Yixuan Zhang, Qianni Luo, Andrea G Parker, and Munmun De Choudhury. 2023. Synthetic lies: Understanding ai-generated misinformation and evaluating algorithmic and human solutions. In *Proceedings of the 2023 CHI conference on human factors in computing systems*. 1–20. doi:10.1145/3544548.3581318