

# Tools for Thought: Research and Design for Understanding, Protecting, and Augmenting Human Cognition with Generative AI

Lev Tankelevitch  
lev.tankelevitch@microsoft.com  
Microsoft Research  
Cambridge, United Kingdom

Majeed Kazemitabaar  
majeed@dgp.toronto.edu  
University of Toronto  
Canada

Srishti Palani  
srishti.palani@salesforce.com  
Tableau Research  
USA

Elena L. Glassman  
glassman@seas.harvard.edu  
Harvard University  
USA

Aniket Kittur  
nkittur@andrew.cmu.edu  
Carnegie Mellon University  
USA

Advait Sarkar  
advait@microsoft.com  
Microsoft Research  
Cambridge, United Kingdom

Jessica He  
jessicahe@ibm.com  
IBM Research  
USA

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

Gonzalo Ramos  
goramos@microsoft.com  
Microsoft Research  
Redmond, USA

Yvonne Rogers  
y.rogers@ucl.ac.uk  
University College London  
London, United Kingdom

Hari Subramonyam  
harihars@stanford.edu  
Stanford University  
USA

## Abstract

We invite researchers, designers, practitioners, and provocateurs to explore what it means to understand and shape the impact of Generative AI (GenAI) on human cognition. GenAI radically widens the scope and capability of automation for work, learning, and creativity. While impactful, it also changes workflows and the quality of thinking involved, raising questions about its effects on cognition, including critical thinking and learning. Yet, GenAI also offers opportunities for designing tools for thought that *protect* and *augment* cognition. Such systems provoke critical thinking, provide personalized tutoring, or enable novel ways of sensemaking, among other approaches. How does GenAI change workflows and human cognition? What are opportunities and challenges for designing GenAI systems that protect and augment human cognition? Which theories, perspectives, and methods are relevant? This workshop aims to develop a multidisciplinary community interested in exploring these questions to protect against the erosion, and fuel the augmentation, of human cognition using GenAI.

## CCS Concepts

• **Human-centered computing** → **Human computer interaction (HCI)**; • **Computing methodologies** → **Artificial intelligence**.

## Keywords

generative AI, artificial intelligence, critical thinking, reasoning, cognition, metacognition, learning, diversity, creativity, sensemaking, autonomy, augmentation, intentionality, reflection, social science, research, design, workshop

## ACM Reference Format:

Lev Tankelevitch, Elena L. Glassman, Jessica He, Majeed Kazemitabaar, Aniket Kittur, Mina Lee, Srishti Palani, Advait Sarkar, Gonzalo Ramos, Yvonne Rogers, and Hari Subramonyam. 2025. Tools for Thought: Research and Design for Understanding, Protecting, and Augmenting Human Cognition with Generative AI. In . ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

## 1 Motivation

“Not another Generative AI workshop!”, you might think to yourself—yet, although Generative AI (GenAI) is an essential aspect of the workshop, our distinctive concern is *human cognition*. This workshop aims to bridge (a) an emerging science of how the use of GenAI affects human thought, from metacognition, to critical thinking, to memory, to creativity (Section 1.1), with (b) an emerging design practice for building GenAI tools that not only *protect* human thought, but also *augment* it (Section 1.2). We aim to gather a multidisciplinary community set on protecting and augmenting human cognition with Generative AI, and kickstart an enduring conversation by mapping the space of opportunities in this area (Section 2).

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](mailto:permissions@acm.org).

Conference'17, July 2017, Washington, DC, USA

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

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM

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

## 1.1 How will people learn and think in the future, as Generative AI is embedded into our work and lives?

The Generative AI shift [12, 64] is upon us. Society is undergoing a radical widening in the scope and capability of automation for work, learning, and creativity. While powerful, there are risks that such automation may bring to human cognition. GenAI systems are commonly deployed as replacements for tools and processes [69]—such as writing and research—through which people think and, therefore, learn, build skills, and grow expertise. GenAI systems also shift workflows from production to ‘critical integration’ [62] of material, involving decisions about when and how to use GenAI, framing tasks, and assessing outputs [35, 55, 63, 72, 76]. Thus, the quality and nature of thinking—where thinking is applied, what kind of thinking is applied—all change when knowledge workflows reorient themselves around GenAI.

GenAI systems therefore act as functional extensions of human cognition, as suggested by the ‘extended mind thesis’ [6].<sup>1</sup> Indeed, GenAI systems and users may form cognitively coupled systems, defined by an interaction between users’ inputs (e.g., fine-tuning, prompting), and systems’ short- and long- term impact on users’ task performance and cognition, with these two components “continuously co-affecting one another in a self-reinforcing cognitive loop” [68]. What changes in human cognition occur within this cognitive loop?

There is evidence that recent technologies are reshaping human cognition. For example, reliance on GPS may negatively impact people’s sense of direction and navigational knowledge [27, 54], and internet use may change the way we process information [52], recall information [70], or assess our confidence in our own knowledge or abilities [24, 52, 61, 71]. Emerging evidence suggests that GenAI may have analogous impacts on cognition. For instance, in education, students with unregulated access to GenAI perform worse when AI is taken away, as they become over-reliant and fail to develop essential independent problem-solving skills [11]. Although directly analogous findings in knowledge work are scant, studies suggest that workers over-rely on these systems in a similar way, at least in the short term, particularly when using GenAI to automate rather than augment tasks [8, 17, 21]. Understanding the nature of this change is vital for making well-informed decisions about protecting aspects of human cognition in light of the GenAI shift, and for designing systems that do so.

Yet the impact we propose to understand is not all negative—well-designed GenAI systems can improve learning [11, 38, 46, 47, 84], writing [87], creativity [7], sensemaking [22, 73], and various forms of knowledge work productivity [15, 20, 21]. Understanding the cognitive underpinnings of this positive impact can help build better AI models and design better systems (e.g., [40]). More broadly, we emphasize that seeking to understand the impact on cognition is complementary to the work on improving AI models or automating tasks. By understanding where, how, and why any impact on our cognition occurs, we can deploy GenAI systems intentionally,

maximizing benefits while mitigating harms. Moreover, by understanding the interplay between GenAI and human cognition, we can design systems that ultimately *augment* it, leading to better short- and long-term outcomes for people.

## 1.2 What opportunities are there for GenAI to augment, or even transform thinking, much as older tools such as writing have done?

While the unconsidered use of GenAI can pose risks to human cognition, GenAI also offers opportunities beyond ‘mere’ task automation—it can help build *tools for thought*, as envisioned by Engelbart [60], Kay [42], and others since then [53, 79]. Alongside automating tasks and generating content, it can help us better understand tasks and think more critically about our work. In addition to increasing the *quantity* of our outputs, it can improve their *quality* by helping us ask better questions, as explored in recent work [16, 19, 36, 65, 67, 74, 86, 87].

One line of work suggests that GenAI systems can support thinking by acting ‘antagonistically’ [16], as ‘provocateurs’ [65], as ‘coaches’ [36], or otherwise prompting ‘metacognitive’ reflection [74], and thereby challenging users’ thinking. Empirical work has begun to explore this direction in individual [32, 57, 58, 83] and collaborative contexts [18, 56]. Relatedly, research is also exploring GenAI as an explicitly educational technology, offering personalized and dynamic tutoring to students [40, 46, 47, 49, 80]. Finally, other work is exploring GenAI as a medium for dynamically representing information at different abstractions [4]. Rather than directly probing users’ thinking, this approach supports cognition by adapting informational representations to the task at hand or encouraging different ways of sensemaking [39, 41, 51, 73, 77, 87] and verification [45, 82]. In sum, this work aims to protect cognition by engaging people in meaningful and beneficial forms of thinking, and moreover, to augment it by challenging or reframing perceptions, assumptions, and understanding. Together with research on understanding the impact of GenAI on cognition, this work offers a pathway for GenAI that not only improves productivity (on top of that afforded by automation), but also leads to improved decision-making, learning, creativity, and understanding.

## 2 Workshop Aims

With GenAI’s rapid real-world adoption and the proliferation of GenAI-driven systems, now is the time to initiate a dedicated effort in understanding what this means for human cognition, and how we can shape this technology for our wide-ranging benefit. The CHI community has both the duty and the means to explore this. As alluded to above, this effort requires multi-disciplinary collaboration. CHI is a unique venue that brings together not only HCI technical system designers, developers, and engineers, but also those from psychology, linguistics, anthropology, sociology, learning science, communications, management, science studies, history, and policy, among others. Similar efforts have arisen in the past, with areas such as Explainable AI [25] and Responsible AI [1, 10] developing strong communities of interest and numerous resources [2, 5, 13, 37]. Cognitive concepts have been alluded to in these contexts, but have not been their core focus—in contrast, our effort is distinctly concerned with the impacts of GenAI on human

<sup>1</sup>Popper’s ‘three worlds’ theory provides a complementary perspective, in which GenAI systems, as products and physical instantiations of ‘World 2’ human cognition and ‘World 3’ ideas, continue to shape our cognition and ideas in both of these worlds [14].

cognition and the implications for the way people learn, build skills, and grow expertise. Thus, this workshop has two aims:

- (1) As the initial instance of an anticipated series, we want to **identify and coalesce a multidisciplinary community** interested in understanding and shaping the impact of GenAI on human cognition.
- (2) We want to **map the space of opportunities** in this area, in terms of the following high-level questions:
  - How do GenAI systems change workflows? What activities become more or less important? What activities are eliminated and created?
  - What is the impact of current GenAI paradigms on human cognition, such as our ability to learn, think critically and creatively, reason, remember, and do sensemaking in contexts from creative art to programming or scientific research? How can we better understand and mitigate potential negative impacts, and maximize positive impacts?
  - What does it mean to protect and augment human cognition?
  - What are opportunities and challenges for designing GenAI systems that protect and augment human cognition? How do we develop principles or guidelines for doing so?
  - Which contemporary and historical theories and perspectives—from across disciplines—are relevant for understanding the impact of GenAI on cognition and workflows, and for designing GenAI systems that protect and augment human cognition?
  - How do we study the impact of current GenAI paradigms and novel GenAI systems for protecting and augmenting human cognition?

Given the breadth and complexity of these questions, we do not expect this initial workshop to provide definitive answers to the above, but rather to initiate focused conversations that lay the groundwork for addressing these questions as a community.

An underlying aim of this workshop and its community is to strengthen the link between *understanding* the impact of GenAI on cognition and *designing* new systems that protect and augment cognition. For example, exploring foundational theories and approaches to studying impact from across relevant disciplines is not only useful for its own sake, but will provide principled grounds for designing novel GenAI systems and evaluating their afforded experiences and effectiveness. The connection between descriptive theories and prescriptive design is illustrated in recent studies on AI-enabled tutoring systems [11, 47], which examine the impact of current GenAI systems on learning, harness pedagogical principles to inform the design of new systems, and conduct comparative evaluations. There are similar opportunities for designing GenAI systems that support critical thinking [65, 76] or sensemaking [19, 73] in knowledge work. CHI's multidisciplinary community, its focus on contextual understanding of technologies, and its strength in designing real-world systems, therefore make it the ideal venue for this workshop.

### 3 Organizers

**Lev Tankelevitch** is a Senior Researcher in Microsoft Research, within the *Tools for Thought* group. His research explores how

to augment human agency in collaborative knowledge work, including using metacognition as a lens to understand and improve human-AI interaction, and to design GenAI systems that improve intentionality in collaboration. He has a background in applied behavioural science, having previously worked at the Behavioural Insights Team, and in cognitive psychology and neuroscience. Find more about his work here: <https://aka.ms/levt>.

**Elena Glassman** is an Assistant Professor at Harvard University's Paulson School Of Engineering and Applied Sciences. Her research focuses on building AI-resilient interfaces [29] that support meta-cognition through a variety of novel interface features and affordances that enhance user's reading, writing, and sensemaking abilities. These features include constructive antagonism [16] and reifying (1) multiple recursive levels of detail [33], (2) natural language commands as dynamically generated UI widgets [78], and (3) pre-computed similarities and distinctions across many items in a corpus at both the literal and analogical levels, e.g., OVER-CODE [30], EXAMPLORE [31], PARALIB [85], and Positional Diction Clustering [28]. Her group specializes in leveraging theories of human cognition about how humans form mental models from varying concrete examples.

**Aniket Kittur** is a Professor in the Human-Computer Interaction Institute in Carnegie Mellon's School of Computer Science. His research explores combining human and machine intelligence to scale up sensemaking and innovation in domains including scientific literature, decision making, productivity, and analogical design.

**Mina Lee** is an Assistant Professor in Computer Science, Data Science Institute, and Cognitive Science (affiliated) at the University of Chicago. Her research centers around Writing with AI, especially how AI is transforming our writing process, the content we produce, and our identities as writers. She was named one of MIT Technology Review's Korean Innovators under 35 in 2022. She has co-founded and organized workshops on Intelligent and Interactive Writing Assistants (In2Writing) and Human-centered Evaluation and Auditing of Language Models (HEAL) at ACL 2022, CHI 2023, and CHI 2024.

**Srishti Palani** is a Senior Researcher at Tableau Research. She researches at the intersection of Cognitive Science, Human-Centered AI and Human-Computer Interaction. Her research investigates how people think and behave while exploring, sensemaking and being creative and with Generative AI and information on the Web. Based on this understanding of user behaviors and cognition, she builds interactive intelligent systems that augment these human cognitive abilities. She earned her doctorate and master's from the University of California, San Diego. Before PhD, she graduated *summa cum laude* double majoring in Computer Science and Psychology (specializing in Cognitive Neuroscience) from Mount Holyoke College.

**Majeed Kazemitabaar** is a PhD candidate at University of Toronto, where he is researching on and developing tools that balance productivity and cognitive engagement in AI-assisted programming. He has studied the impact of learning to code with AI on subsequent performance without AI, to measure the effects of overreliance on AI [43, 44]. He has developed and evaluated programming tools based on the concept of "*Friction-Induced AI*" to achieve two goals: enhancing short-term productivity by improving verification through added intervention points [45], and preventing

long-term productivity loss by requiring user engagement with AI-generated code [46].

**Jessica He** is a UX Designer at IBM Research, where she is a member of the Human-AI Collaboration team. Her work focuses on leveraging design to bridge the gap between user expectations and emerging AI technologies, encompassing topics including AI attribution, risk mitigation, and enhancing knowledge work [34]. By applying user-centered methods, she strives to create - and guide other practitioners in creating - trustworthy generative AI applications that augment rather than replace human capabilities and collaboration [81].

**Gonzalo Ramos** is a Principal Researcher at Microsoft Research at Redmond. He is part of the Human Centered AI and Experiences Group at Microsoft Research at Redmond, where he works at the intersection of HCI, Design, and AI to augment people's agencies and capabilities. He is a graduate from the University of Toronto's DGP lab, as well as the Universidad de Buenos Aires. Prior to his position at Microsoft Research, he worked as a Senior Design Technologist and later UX Scientist at Amazon, and as a Scientist at Microsoft. You can find more information on his work here: <https://www.microsoft.com/en-us/research/people/goramos/>.

**Advait Sarkar** is a researcher at Microsoft, affiliated lecturer at the University of Cambridge, and honorary lecturer at University College London. He studies the effects of Generative AI on knowledge work [62, 75], programming [48, 64], and data analysis [23, 26, 45, 50]. He is part of the *Tools for Thought* group at Microsoft [59], where he leads a research agenda aimed at enhancing critical thinking with Generative AI [66, 67]. His article "*AI Should Challenge, Not Obey*" appears in the October 2024 issue of Communications of the ACM [65].

**Yvonne Rogers** is a Professor of Interaction Design at University College London. A central theme of her work is concerned with designing AI that augments human cognition. She has given many keynotes and invited talks on how HCI can meet AI in the I.

**Hari Subramonyam** is an Assistant Professor (Research) at Stanford Graduate School of Education and Computer Science (by courtesy). He is also the Ram and Vijay Sriram Faculty Fellow at the Stanford Institute for Human-Centered AI. Subramonyam studies ways to augment human learning using AI by (1) engaging in cognitively informed design practices, (2) co-designing with learners and educators, and (3) developing transformative AI-enabled learning experiences. Through his research, he also contributes tools and methodologies that prioritize ethical considerations, responsible design practices, and human values when creating AI experiences.

## 4 Workshop Organisation: In Person

The workshop will be in-person only to facilitate in-depth discussion among all participants, and to avoid challenges in ensuring equitable opportunities for participation between in-person and remote attendees. While we acknowledge that this decision would unfortunately exclude those who cannot travel to attend in person, we believe that this trade-off would result in a superior experience for those who do attend, and is better aligned with the primary community-building aim of the workshop. Alongside synchronous

in-person engagement, we will use Discord for asynchronous discussions before, during, and after the workshop. To facilitate in-depth discussion and community-building, we are designing the workshop for 30 participants. Presentations and other media will be shown on the provided projector and also shared on the asynchronous platform. The workshop is planned as a 1-day event, 9:00AM to 5:30PM, with breaks.

### 4.1 Accessibility

Authors will be encouraged to submit position papers or other media in accessible versions. We will use a platform that enables captions for speech. We will also determine participants' accessibility needs in advance and will liaise with the CHI Accessibility Chair if further support is needed.

### 4.2 Asynchronous Materials

We plan to use a dedicated website (<https://ai-tools-for-thought.github.io/workshop/>) to promote and distribute workshop information, including the call for proposals, publish workshop submissions in advance of the workshop, and publish a summary of activities after the workshop. We will create a Discord channel for asynchronous communication and media sharing before, during, and after the workshop.

## 5 Workshop Activities

### 5.1 Before the Workshop

We aim to build a multidisciplinary community and therefore hope to attract participants not only from HCI, including researchers, designers, technologists, and industry professionals, but also those studying and working in related fields such as psychology, education, data science, anthropology, sociology, philosophy, and others. We will advertise the workshop through its website, social media (e.g., X/Twitter, LinkedIn), email distribution lists from relevant conferences (including CHI, CSCW, IUI, ACL) and research and industry institutions, and direct communication with colleagues. The diversity of organisers' institutions should facilitate this dissemination.

Those wishing to participate will be asked to submit either a position paper (of up to 2,000 words) or other portfolio of work (e.g., a system or set of design explorations) accompanied by a description of its relevance. We particularly welcome submissions that offer novel perspectives, critical reflections, and/or preliminary results on the topic of GenAI's potential to impact cognition or enable the design of tools for thought (see call for proposals in Section 7). Submissions will be reviewed by at least two members of the workshop's organizing committee (corresponding to ACM's 'Reviewed' level of prepublication evaluation [9]). To help increase the diversity of perspectives at the workshop, structure the workshop activities, and facilitate community-building, we will also ask participants to submit a short personal statement of up to 150 words describing their relevant background, their motivation for attending, and specific areas of interest. This will be considered holistically with the submitted work.

Accepted submissions will be posted on the workshop's website in advance of CHI2025, to encourage participants to gain familiarity with each other's work. We will also explore using GenAI

to creatively synthesize and communicate submitted information to provide an overview of the incoming participant community in advance of the workshop, and potentially support workshop activities.

## 5.2 During the Workshop

We will start the workshop with a welcome session outlining the main themes and upcoming sessions. We will also engage participants in an ice-breaker activity: the 'human spectrogram' [3], where participants indicate their opinion about provocative questions central to the workshop by physically moving along an imagined spectrum covering the response extremes (e.g., 'strongly agree' to 'strongly disagree'). This will introduce the range of questions central to the workshop, help participants discover similarities and differences in their opinions, and stimulate discussions.

The workshop will then be divided into four sessions, each lasting 75 minutes, and focusing on different themes (exact themes will be finalized based on the submissions and participants; tentative schedule in Table 1):

- **Session 1 (talks and discussion): Impact of GenAI on Cognition and Workflows.** What is known about the positive and negative impact of GenAI on cognition and workflows? What important perspectives (e.g., historical, cultural, etc.) on these impacts do we need to be aware of?
- **Session 2 (talks and discussion): Methods and Theories.** What are existing (and novel) methods and theories from across disciplines relevant for studying the impact of GenAI on cognition and workflows, and designing and evaluating new GenAI systems? How is HCI research well-positioned (or ill-positioned) to apply these methods and theories?
- **Session 3 (talks and discussion): Designing for Cognitive Protection and Augmentation.** What does it mean to protect and augment cognition? What are promising prototypes, frameworks, design approaches, guidelines, or perspectives for designing GenAI systems that protect and augment cognition?
- **Session 4 (co-ideation): Mapping opportunities.** A co-ideation activity with all attendees in groups, with each group focusing on:
  - Key research questions for understanding the impact of GenAI on cognition and workflows
  - Specific areas of cognition and aspects of workflows that may be impacted by GenAI
  - Opportunities for designing GenAI systems that protect and augment cognition
  - Challenges for designing GenAI systems that protect and augment cognition
  - Potential principles and guidelines for designing GenAI systems that protect and augment cognition

The first three sessions will feature 5-6 lightning talks (each presentation approximately 5 minutes), followed by a panel discussion with all presenters together with audience Q&A (approximately 45 minutes). These sessions will be moderated by a member of the organizing committee. Alongside the live questions, we will encourage participants to submit thoughts, ideas, and questions to Discord. Presenters will be a subset of all participants, selected

by the organizing committee to ensure a diversity of perspectives, approaches, and types of contributions. As our workshop aims to facilitate community-building and mapping of opportunities for the future, this format will ensure ample time for questions and discussion in a way that is accessible to all attendees, while balancing the need to provide presentation opportunities for as many participants as possible.

The fourth session will consist of a co-ideation activity with all attendees, split into groups of 5-6, and allocated based on their indicated interests and submissions. The co-ideation activity will be 50 minutes, followed by 25 minutes of group share-outs. Each group will be provided with discussion prompts, and moderated by at least one workshop organizer.

Coffee breaks will be used to informally showcase any demos or related work that participants have submitted; interested participants will be instructed in advance to prepare their demos. We will organize an informal outing for lunch to facilitate further networking and discussion. We will encourage the use of the asynchronous platform throughout the workshop to facilitate ideation, information exchange, and networking.

The workshop will conclude by summarizing the day's insights, eliciting participants' opinions on next steps for the community, and suggesting future action items.

## 6 Publishing Workshop Proceedings and Post-Workshop Plans

We will publish a collection of the submitted papers as workshop proceedings via CEUR-WS. We will use the workshop website to create a permanent record that links to these archival versions of all position papers, as well as any portfolios, and recordings of all talks and panel sessions.

During the closing segment of the workshop, we will discuss and gather feedback about strategies for disseminating the materials covered during the workshop, such as research questions, positions, and design principles. We anticipate that the discussions from this workshop will inspire an expanded position paper, a collection of articles for special issues of peer-reviewed journals like TOCHI, or a series of articles on platforms like Medium. All output will be posted on the workshop website, and shared via social media.

Furthermore, we will deliberate on methods to cultivate the Tools for Thought community through social media and/or a special interest group (SIG). We will also consider organizing follow-up events in HCI, design, and ML venues to expand the size and visibility of the research community.

## 7 Call for Participation

We invite researchers, designers, and provocateurs to explore what it means to understand and shape the impact of Generative AI (GenAI) on human cognition. GenAI radically widens the scope and capability of automation for work, learning, and creativity. However, it also changes workflows and the quality of thinking involved, raising questions about its impact on human cognition. Nonetheless, GenAI offers opportunities for designing tools for thought—systems that *protect*, and moreover *augment* cognition.

How does GenAI change workflows and impact human cognition, such as our ability to learn, think critically and creatively,

**Table 1: Proposed Workshop Schedule**

| Slot                  | Activity                                                                                       |
|-----------------------|------------------------------------------------------------------------------------------------|
| 09:00 - 09:30 (30min) | Welcome, ice-breaker, and coffee                                                               |
| 09:30 - 10:45 (75min) | Lightning talks + panel discussion: <i>Impact of GenAI on Cognition and Workflows</i>          |
| 10:45 - 11:15 (30min) | Coffee break and informal demos                                                                |
| 11:15 - 12:30 (75min) | Lightning talks + panel discussion: <i>Methods and Theories</i>                                |
| 12:30 - 14:00 (90min) | Lunch break (informal networking and discussion)                                               |
| 14:00 - 15:15 (75min) | Lightning talks + panel discussion: <i>Designing for Cognitive Protection and Augmentation</i> |
| 15:15 - 15:45 (30min) | Coffee break and informal demos                                                                |
| 15:45 - 17:00 (75min) | Co-ideation session: <i>Mapping Opportunities</i>                                              |
| 17:00 - 17:30 (30min) | Next steps and closing                                                                         |

reason, remember, and carry out sensemaking? What are opportunities and challenges for designing GenAI systems that protect and augment human cognition? Which theories, perspectives, and methods are relevant? Through lightning talks, panel discussions, and co-ideation sessions, **this workshop aims to build a multi-disciplinary community interested in these questions** and in mapping the opportunities for addressing them.

We invite position papers (up to 2,000 words) or portfolios of work (accompanied by description of relevance) addressing the above or related questions. We particularly welcome submissions that offer **novel perspectives, critical reflections, revisitations of older yet relevant work, and/or preliminary findings**. Submissions should be in CHI Extended Abstract format, submitted through the workshop website: <https://ai-tools-for-thought.github.io/workshop/>, and will be reviewed by the organizing committee. Accepted submissions will be published as workshop proceedings, posted on the workshop website, and shared on social media, with a subset also selected for a lightning talk and panel discussion. Participants are also asked to submit a short personal statement of relevant background and interests for the workshop (up to 150 words). Submitted work and statements will be considered holistically to increase the diversity of perspectives, structure workshop activities, and facilitate community-building. At least one author of each accepted submission must attend the workshop and all participants must register for both the workshop and at least one day of the conference.

## 8 Expected size of attendance

We plan for 30 in-person attendees only.

## References

- [1] [n. d.]. ACM Conference on Fairness, Accountability, and Transparency (FAccT). <https://faccconference.org/>. Accessed: 2024-10-10.
- [2] [n. d.]. HAX Toolkit - Microsoft Research. <https://www.microsoft.com/en-us/research/project/hax-toolkit/>. Accessed: 2024-10-10.
- [3] [n. d.]. Human Spectrogram - Open Design Kit. <http://opendesignkit.org/methods/human-spectrogram/>. Accessed: 2024-10-10.
- [4] 2024. Dynamic Abstractions Workshop: Building the Next Generation of Cognitive Tools and Interfaces. UIST'24 Workshop, Pittsburgh, USA, October 13-16, 2024. <https://dynamic-uist24.vercel.app/>. Organized by Sang Ho Suh, Hai Dang, Ryan Yen, Josh Pollock, Ian Arawjo, Rubaiat Habib, Nazmus Saquib, Hari Subramonyam, Jingyi, and Arvind Satyanarayan.
- [5] Saleema Amershi, Dan Weld, Mihaela Vorvoreanu, Adam Fourney, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, Paul N Bennett, Kori Inkpen, et al. 2019. Guidelines for human-AI interaction. In *Proceedings of the 2019 chi conference on human factors in computing systems*. 1–13.
- [6] Clark Andy and Chalmers David. 1998. The extended mind. *Analysis* 58, 1 (1998), 7–19.
- [7] Tyler Angert, Miroslav Suzara, Jenny Han, Christopher Pondoc, and Hariharan Subramonyam. 2023. Spellburst: A node-based interface for exploratory creative coding with natural language prompts. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. 1–22.
- [8] Zahra Ashktorab, Qian Pan, Werner Geyer, Michael Desmond, Marina Danilevsky, James M Johnson, Casey Dugan, and Michelle Bachman. 2024. Emerging Reliance Behaviors in Human-AI Text Generation: Hallucinations, Data Quality Assessment, and Cognitive Forcing Functions. *arXiv preprint arXiv:2409.08937* (2024).
- [9] Association for Computing Machinery. 2024. Pre-Publication Evaluation. <https://www.acm.org/publications/policies/pre-publication-evaluation>. Accessed: 2024-10-09.
- [10] Gagan Bansal, Alison Marie Smith-Renner, Zana Bućinca, Tongshuang Wu, Kenneth Holstein, Jessica Hullman, and Simone Stumpf. 2022. Workshop on Trust and Reliance in AI-Human Teams (TRAIT). In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (CHI EA '22). Association for Computing Machinery, New York, NY, USA, Article 116, 6 pages. <https://doi.org/10.1145/3491101.3503704>
- [11] Hamsa Bastani, Osbert Bastani, Alp Sungu, Haosen Ge, Ozge Kabakci, and Rei Mariman. 2024. Generative ai can harm learning. *Available at SSRN* 4895486 (2024).
- [12] Alexander Bick, Adam Blandin, and David Deming. 2024. *The Rapid Adoption of Generative AI*. Technical Report.
- [13] Sarah Bird, Miro Dudik, Richard Edgar, Brandon Horn, Roman Lutz, Vanessa Milan, Mehrnoosh Sameki, Hanna Wallach, and Kathleen Walker. 2020. *Fairlearn: A toolkit for assessing and improving fairness in AI*. Technical Report MSR-TR-2020-32. Microsoft. <https://www.microsoft.com/en-us/research/publication/fairlearn-a-toolkit-for-assessing-and-improving-fairness-in-ai/>
- [14] Brian Boyd. 2016. Popper's world 3: Origins, progress, and import. *Philosophy of the Social Sciences* 46, 3 (2016), 221–241.
- [15] Erik Brynjolfsson, Danielle Li, and Lindsey R. Raymond. 2023. Generative AI at Work. <https://doi.org/10.3386/w31161>
- [16] Alice Cai, Ian Arawjo, and Elena L. Glassman. 2024. Antagonistic AI. <https://arxiv.org/abs/2402.07350v1>
- [17] Zenan Chen and Jason Chan. 2023. Large Language Model in Creative Work: The Role of Collaboration Modality and User Expertise. <https://doi.org/10.2139/ssrn.4575598>
- [18] Chun-Wei Chiang, Zhuoran Lu, Zhuoyan Li, and Ming Yin. 2024. Enhancing AI-Assisted Group Decision Making through LLM-Powered Devil's Advocate. In *Proceedings of the 29th International Conference on Intelligent User Interfaces (IUI '24)*. Association for Computing Machinery, New York, NY, USA, 103–119. <https://doi.org/10.1145/3640543.3645199>
- [19] Katherine M. Collins, Ilia Sucholutsky, Umang Bhatt, Kartik Chandra, Lionel Wong, Mina Lee, Cedegao E. Zhang, Tan Zhi-Xuan, Mark Ho, Vikash Mansinghka, Adrian Weller, Joshua B. Tenenbaum, and Thomas L. Griffiths. 2024. Building Machines that Learn and Think with People. <http://arxiv.org/abs/2408.03943> [cs].
- [20] Zheyuan Kevin Cui, Mert Demirer, Sonia Jaffe, Leon Musolf, Sida Peng, and Tobias Salz. 2024. The Effects of Generative AI on High Skilled Work: Evidence from Three Field Experiments with Software Developers. *Available at SSRN* (2024).
- [21] Fabrizio Dell'Acqua, Edward McFowland, Ethan R Mollick, Hila Lifshitz-Assaf, Katherine Kellogg, Saran Rajendran, Lisa Kray, François Candelon, and Karim R Lakhani. 2023. Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality. *Harvard Business School Technology & Operations Mgt. Unit Working Paper* 24-013

- (2023).
- [22] Ian Drosos, Advait Sarkar, Xiaotong Xu, Carina Negreanu, Sean Rintel, and Lev Tankelevitch. 2024. "It's like a rubber duck that talks back": Understanding Generative AI-Assisted Data Analysis Workflows through a Participatory Prompting Study. In *Proceedings of the 3rd Annual Meeting of the Symposium on Human-Computer Interaction for Work*. ACM, Newcastle upon Tyne United Kingdom, 1–21. <https://doi.org/10.1145/3663384.3663389>
  - [23] Ian Drosos, Advait Sarkar, Xiaotong Xu, Carina Negreanu, Sean Rintel, and Lev Tankelevitch. 2024. "It's like a rubber duck that talks back": Understanding Generative AI-Assisted Data Analysis Workflows through a Participatory Prompting Study. In *Proceedings of the 3rd Annual Meeting of the Symposium on Human-Computer Interaction for Work* (Newcastle upon Tyne, United Kingdom) (*CHIWORK '24*). Association for Computing Machinery, New York, NY, USA, Article 16, 21 pages. <https://doi.org/10.1145/3663384.3663389>
  - [24] Timothy L. Dunn, Connor Gaspar, Daev McLean, Derek J. Koehler, and Evan F. Risko. 2021. Distributed metacognition: Increased bias and deficits in metacognitive sensitivity when retrieving information from the internet. *Technology, Mind, and Behavior* 2, 3 (Aug. 2021). <https://doi.org/10.1037/tmb0000039>
  - [25] Upol Ehsan, Philipp Wintersberger, Q Vera Liao, Martina Mara, Marc Streit, Sandra Wachter, Andreas Riener, and Mark O Riedl. 2021. Operationalizing human-centered perspectives in explainable AI. In *Extended abstracts of the 2021 CHI conference on human factors in computing systems*. 1–6.
  - [26] Kasra Ferdowsi, Jack Williams, Ian Drosos, Andrew D. Gordon, Carina Negreanu, Nadia Polikarpova, Advait Sarkar, and Benjamin Zorn. 2023. COLDECO: An End User Spreadsheet Inspection Tool for AI-Generated Code. In *2023 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*. 82–91. <https://doi.org/10.1109/VL-HCC57772.2023.00017>
  - [27] Pablo Fernandez-Velasco and Hugo J Spiers. 2024. Wayfinding across ocean and tundra: what traditional cultures teach us about navigation. *Trends in Cognitive Sciences* 28, 1 (2024), 56–71.
  - [28] Katy Ilonka Gero, Chelse Swoopes, Ziwei Gu, Jonathan K. Kummerfeld, and Elena L. Glassman. 2024. Supporting Sensemaking of Large Language Model Outputs at Scale. 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 838, 21 pages. <https://doi.org/10.1145/3613904.3642139>
  - [29] Elena L. Glassman, Ziwei Gu, and Jonathan K Kummerfeld. 2024. AI-Resilient Interfaces. *arXiv preprint arXiv:2405.08447* (2024).
  - [30] Elena L. Glassman, Jeremy Scott, Rishabh Singh, Philip J Guo, and Robert C Miller. 2015. OverCode: Visualizing variation in student solutions to programming problems at scale. *ACM Transactions on Computer-Human Interaction (TOCHI)* 22, 2 (2015), 1–35.
  - [31] Elena L. Glassman, Tianyi Zhang, Björn Hartmann, and Miryung Kim. 2018. Visualizing API usage examples at scale. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–12.
  - [32] Frederic Gmeiner, Humphrey Yang, Lining Yao, Kenneth Holstein, and Nikolas Martelaro. 2023. Exploring Challenges and Opportunities to Support Designers in Learning to Co-create with AI-based Manufacturing Design Tools. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, 1–20. <https://doi.org/10.1145/3544548.3580999>
  - [33] Ziwei Gu, Ian Arawjo, Kenneth Li, Jonathan K Kummerfeld, and Elena L. Glassman. 2024. An AI-Resilient Text Rendering Technique for Reading and Skimming Documents. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 1–22.
  - [34] Jessica He, Stephanie Houde, Gabriel E Gonzalez, Dario Andrés Silva Moran, Steven I Ross, Michael Muller, and Justin D Weisz. 2024. AI and the Future of Collaborative Work: Group Ideation with an LLM in a Virtual Canvas. In *Proceedings of the 3rd Annual Meeting of the Symposium on Human-Computer Interaction for Work*. 1–14.
  - [35] Jessica He, David Piorkowski, Michael Muller, Kristina Brimijoin, Stephanie Houde, and Justin Weisz. 2023. Rebalancing worker initiative and AI initiative in future work: Four task dimensions. In *Proceedings of the 2nd Annual Meeting of the Symposium on Human-Computer Interaction for Work*. 1–16.
  - [36] Jake M. Hofman, Daniel G. Goldstein, and David M. Rothschild. 2023. A Sports Analogy for Understanding Different Ways to Use AI. *Harvard Business Review* (Dec. 2023). <https://hbr.org/2023/12/a-sports-analogy-for-understanding-different-ways-to-use-ai> Section: AI and machine learning.
  - [37] Matthew K Hong, Adam Fournery, Derek DeBellis, and Saleema Amershi. 2021. Planning for natural language failures with the ai playbook. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–11.
  - [38] Xinying Hou, Zihan Wu, Xu Wang, and Barbara J. Ericson. 2024. CodeTailor: LLM-Powered Personalized Parsons Puzzles for Engaging Support While Learning Programming. In *Proceedings of the Eleventh ACM Conference on Learning @ Scale* (Atlanta, GA, USA) (*L@S '24*). Association for Computing Machinery, New York, NY, USA, 51–62. <https://doi.org/10.1145/3657604.3662032>
  - [39] Peiling Jiang, Jude Rayan, Steven P. Dow, and Haijun Xia. 2023. Graphologue: Exploring Large Language Model Responses with Interactive Diagrams. <https://doi.org/10.1145/3586183.3606737> arXiv:2305.11473 [cs].
  - [40] Irina Jurenka, Markus Kunesch, Kevin R McKee, Daniel Gillick, Shaojian Zhu, Sara Wiltberger, Shubham Milind Phal, Katherine Hermann, Daniel Kasenberg, Avishkar Bhoopchand, et al. 2024. Towards responsible development of generative AI for education: An evaluation-driven approach. *arXiv preprint arXiv:2407.12687* (2024).
  - [41] Hyeonsu B Kang, Tongshuang Wu, Joseph Chee Chang, and Aniket Kittur. 2023. Synergi: A Mixed-Initiative System for Scholarly Synthesis and Sensemaking. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. 1–19.
  - [42] Alan Kay. 1990. User Interface: A Personal View. In *The Art of Human-Computer Interface Design*. 191–207. [http://ui.korea.ac.kr/Board/Upload/a%20personal%20view\\_n.pdf](http://ui.korea.ac.kr/Board/Upload/a%20personal%20view_n.pdf)
  - [43] Majeed Kazemitabaar, Justin Chow, Carl Ka To Ma, Barbara J. Ericson, David Weintrop, and Tovi Grossman. 2023. Studying the effect of AI Code Generators on Supporting Novice Learners in Introductory Programming. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, 1–23. <https://doi.org/10.1145/3544548.3580919>
  - [44] Majeed Kazemitabaar, Xinying Hou, Austin Henley, Barbara J. Ericson, David Weintrop, and Tovi Grossman. 2023. How Novices Use LLM-Based Code Generators to Solve CS1 Coding Tasks in a Self-Paced Learning Environment. <https://doi.org/10.48550/arXiv.2309.14049> arXiv:2309.14049 [cs].
  - [45] Majeed Kazemitabaar, Jack Williams, Ian Drosos, Tovi Grossman, Austin Henley, Carina Negreanu, and Advait Sarkar. 2024. Improving Steering and Verification in AI-Assisted Data Analysis with Interactive Task Decomposition. *arXiv preprint arXiv:2407.02651* (2024).
  - [46] Majeed Kazemitabaar, Runlong Ye, Xiaoning Wang, Austin Zachary Henley, Paul Denny, Michelle Craig, and Tovi Grossman. 2024. Codeaid: Evaluating a classroom deployment of an llm-based programming assistant that balances student and educator needs. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 1–20.
  - [47] Harsh Kumar, David M. Rothschild, Daniel G. Goldstein, and Jake M. Hofman. 2023. Math Education with Large Language Models: Peril or Promise? <https://doi.org/10.2139/ssrn.4641653>
  - [48] Michael Jing Long Lee, Advait Sarkar, and Alan F. Blackwell. 2024. Predictability of identifier naming with Copilot: A case study for mixed-initiative programming tools. In *Proceedings of the 35th Annual Conference of the Psychology of Programming Interest Group (PPIG 2024)*.
  - [49] Wengxi Li, Roy Pea, Nick Haber, and Hari Subramonyam. 2024. Tutorly: Turning Programming Videos Into Apprenticeship Learning Environments with LLMs. *arXiv preprint arXiv:2405.12946* (2024).
  - [50] Michael Xieyang Liu, Advait Sarkar, Carina Negreanu, Benjamin Zorn, Jack Williams, Neil Toronto, and Andrew D. Gordon. 2023. "What It Wants Me To Say": Bridging the Abstraction Gap Between End-User Programmers and Code-Generating Large Language Models. 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 598, 31 pages. <https://doi.org/10.1145/3544548.3580817>
  - [51] Michael Xieyang Liu, Tongshuang Wu, Tianying Chen, Franklin Mingzhe Li, Aniket Kittur, and Brad A Myers. 2024. Selenite: Scaffolding Online Sensemaking with Comprehensive Overviews Elicited from Large Language Models. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 1–26.
  - [52] Elizabeth J. Marsh and Suparna Rajaram. 2019. The Digital Expansion of the Mind: Implications of Internet Usage for Memory and Cognition. *Journal of Applied Research in Memory and Cognition* 8, 1 (March 2019), 1–14. <https://doi.org/10.1016/j.jarmac.2018.11.001>
  - [53] Andy Matuschak, Michael Nielsen, Andy Matuschak, and Michael Nielsen. 2019. How can we develop transformative tools for thought? *Numinous. Productions (blog)* (12.01. 2020) (2019).
  - [54] Laura Miola, Veronica Muffato, Enrico Sella, Chiara Meneghetti, and Francesca Pazzaglia. 2024. GPS Use and Navigation Ability: A Systematic Review and Meta-Analysis. *Journal of Environmental Psychology* (2024), 102417.
  - [55] Srishti Palani and Gonzalo Ramos. 2024. Evolving Roles and Workflows of Creative Practitioners in the Age of Generative AI. In *Proceedings of the 16th Conference on Creativity & Cognition*. 170–184.
  - [56] Gun Woo (Warren) Park, Payod Panda, Lev Tankelevitch, and Sean Rintel. 2024. The CoExplorer Technology Probe: A Generative AI-Powered Adaptive Interface to Support Intentionality in Planning and Running Video Meetings. In *Designing Interactive Systems Conference (DIS '24, Vol. 4)*. ACM, 1638–1657. <https://doi.org/10.1145/3643834.3661507>
  - [57] Soya Park and Chinmay Kulkarni. 2024. Thinking Assistants: LLM-Based Conversational Assistants that Help Users Think By Asking rather than Answering. <http://arxiv.org/abs/2312.06024> arXiv:2312.06024 [cs].
  - [58] Leon Reicherts, Gun Woo Park, and Yvonne Rogers. 2022. Extending Chatbots to Probe Users: Enhancing Complex Decision-Making Through Probing Conversations. In *Proceedings of the 4th Conference on Conversational User Interfaces (CUI '22)*. Association for Computing Machinery, New York, NY, USA, 1–10.

- <https://doi.org/10.1145/3543829.3543832>
- [59] Microsoft Research. 2024. Tools for Thought. <https://www.microsoft.com/en-us/research/project/t4t/>. Accessed: 2024-09-26.
  - [60] Howard Rheingold. 1985. *Tools for Thought: The People and Ideas behind the Next Computer Revolution*. Simon & Schuster Trade.
  - [61] Evan F. Risko, Amanda M. Ferguson, and David McLean. 2016. On retrieving information from external knowledge stores: Feeling-of-findability, feeling-of-knowing and Internet search. *Computers in Human Behavior* 65 (Dec. 2016), 534–543. <https://doi.org/10.1016/j.chb.2016.08.046>
  - [62] Advait Sarkar. 2023. Exploring Perspectives on the Impact of Artificial Intelligence on the Creativity of Knowledge Work: Beyond Mechanised Plagiarism and Stochastic Parrots. In *Proceedings of the 2nd Annual Meeting of the Symposium on Human-Computer Interaction for Work* (Oldenburg, Germany) (CHIWORK '23). Association for Computing Machinery, New York, NY, USA, Article 13, 17 pages. <https://doi.org/10.1145/3596671.3597650>
  - [63] Advait Sarkar. 2023. Exploring Perspectives on the Impact of Artificial Intelligence on the Creativity of Knowledge Work: Beyond Mechanised Plagiarism and Stochastic Parrots. In *Proceedings of the 2nd Annual Meeting of the Symposium on Human-Computer Interaction for Work* (Oldenburg, Germany) (CHIWORK '23). Association for Computing Machinery, New York, NY, USA, Article 13, 17 pages. <https://doi.org/10.1145/3596671.3597650>
  - [64] Advait Sarkar. 2023. Will Code Remain a Relevant User Interface for End-User Programming with Generative AI Models?. In *Proceedings of the 2023 ACM SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software* (Cascais, Portugal) (Onward! 2023). Association for Computing Machinery, New York, NY, USA, 153–167. <https://doi.org/10.1145/3622758.3622882>
  - [65] Advait Sarkar. 2024. AI Should Challenge, Not Obey. *Commun. ACM* (Sept. 2024), 5 pages. <https://doi.org/10.1145/3649404> Online First.
  - [66] Advait Sarkar. 2024. Intention Is All You Need. *Proceedings of the 35th Annual Conference of the Psychology of Programming Interest Group (PPIG 2024)* (2024).
  - [67] Advait Sarkar, Xiaotong (Tone) Xu, Neil Toronto, Ian Drosos, and Christian Poelitz. 2024. When Copilot Becomes Autopilot: Generative AI's Critical Risk to Knowledge Work and a Critical Solution. In *Proceedings of the Annual Conference of the European Spreadsheet Risks Interest Group (EuSpRIG 2024)*.
  - [68] Wout Schellaert, Fernando Martínez-Plumed, Karina Vold, John Burden, Pablo A. M. Casares, Bao Sheng Loe, Roi Reichart, Sean Ó hÉigeartaigh, Anna Korhonen, and José Hernández-Orallo. 2023. Your Prompt is My Command: On Assessing the Human-Centred Generality of Multimodal Models. *Journal of Artificial Intelligence Research* 77 (June 2023), 377–394. <https://doi.org/10.1613/jair.1.14157>
  - [69] Auste Simkute, Lev Tankelevitch, Viktor Kewenig, Ava Elizabeth Scott, Abigail Sellen, and Sean Rintel. 2024. Ironies of Generative AI: Understanding and mitigating productivity loss in human-AI interactions. <https://doi.org/10.48550/arXiv.2402.11364> arXiv:2402.11364 [cs].
  - [70] Betsy Sparrow, Jenny Liu, and Daniel M Wegner. 2011. Google effects on memory: Cognitive consequences of having information at our fingertips. *science* 333, 6043 (2011), 776–778.
  - [71] Sean M. Stone and Benjamin C. Storm. 2021. Search fluency as a misleading measure of memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 47, 1 (2021), 53–64. <https://doi.org/10.1037/xlm0000806> Place: US Publisher: American Psychological Association.
  - [72] Hariharan Subramonyam, Roy Pea, Christopher Lawrence Pondoc, Maneesh Agrawala, and Colleen Seifert. 2024. Bridging the Gulf of Envisioning: Cognitive Design Challenges in LLM Interfaces. <https://doi.org/10.48550/arXiv.2309.14459> arXiv:2309.14459 [cs].
  - [73] Sangho Suh, Bryan Min, Srishti Palani, and Haijun Xia. 2023. Sensecape: Enabling Multilevel Exploration and Sensemaking with Large Language Models. <https://doi.org/10.48550/arXiv.2305.11483> arXiv:2305.11483 [cs].
  - [74] Lev Tankelevitch, Viktor Kewenig, Auste Simkute, Ava Elizabeth Scott, Advait Sarkar, Abigail Sellen, and Sean Rintel. 2024. The Metacognitive Demands and Opportunities of Generative AI. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA, 1–24. <https://doi.org/10.1145/3613904.3642902>
  - [75] Lev Tankelevitch, Viktor Kewenig, Auste Simkute, Ava Elizabeth Scott, Advait Sarkar, Abigail Sellen, and Sean Rintel. 2024. The Metacognitive Demands and Opportunities of Generative AI. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (Honolulu, HI, USA) (CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 680, 24 pages. <https://doi.org/10.1145/3613904.3642902>
  - [76] Lev Tankelevitch, Viktor Kewenig, Auste Simkute, Ava Elizabeth Scott, Advait Sarkar, Abigail Sellen, and Sean Rintel. 2024. The Metacognitive Demands and Opportunities of Generative AI. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (Honolulu, HI, USA) (CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 680, 24 pages. <https://doi.org/10.1145/3613904.3642902>
  - [77] Priyan Vaithilingam, Ian Arawjo, and Elena L Glassman. 2024. Imagining a future of designing with ai: Dynamic grounding, constructive negotiation, and sustainable motivation. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference*. 289–300.
  - [78] Priyan Vaithilingam, Elena L Glassman, Jeevana Priya Inala, and Chenglong Wang. 2024. DynaVis: Dynamically Synthesized UI Widgets for Visualization Editing. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 1–17.
  - [79] Bret Victor. 2014. Humane representation of thought: a trail map for the 21st century. In *Proceedings of the companion publication of the 2014 ACM SIGPLAN conference on Systems, Programming, and Applications: Software for Humanity*. 5–5.
  - [80] Rose Wang, Qingyang Zhang, Carly Robinson, Susanna Loeb, and Dorottya Demszky. 2024. Bridging the Novice-Expert Gap via Models of Decision-Making: A Case Study on Remediating Math Mistakes. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, Kevin Duh, Helena Gomez, and Steven Bethard (Eds.). Association for Computational Linguistics, Mexico City, Mexico, 2174–2199. <https://doi.org/10.18653/v1/2024.naacl-long.120>
  - [81] Justin D Weisz, Jessica He, Michael Muller, Gabriela Hofer, Rachel Miles, and Werner Geyer. 2024. Design Principles for Generative AI Applications. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 1–22.
  - [82] Liwenhan Xie, Chengbo Zheng, Haijun Xia, Huamin Qu, and Chen Zhu-Tian. 2024. WaitGPT: Monitoring and Steering Conversational LLM Agent in Data Analysis with On-the-Fly Code Visualization. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology* (Pittsburgh, PA, USA) (UIST '24). Association for Computing Machinery, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3654777.3676374>
  - [83] Xiaotong (Tone) Xu, Jiayu Yin, Catherine Gu, Jenny Mar, Sydney Zhang, Jane L. E, and Steven P. Dow. 2024. Jamplate: Exploring LLM-Enhanced Templates for Idea Reflection. In *Proceedings of the 29th International Conference on Intelligent User Interfaces (IUI '24)*. Association for Computing Machinery, New York, NY, USA, 907–921. <https://doi.org/10.1145/3640543.3645196>
  - [84] Litao Yan, Alyssa Hwang, Zhiyuan Wu, and Andrew Head. 2024. Ivie: Lightweight Anchored Explanations of Just-Generated Code. 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 140, 15 pages. <https://doi.org/10.1145/3613904.3642239>
  - [85] Litao Yan, Miryung Kim, Bjoern Hartmann, Tianyi Zhang, and Elena L Glassman. 2022. Concept-annotated examples for library comparison. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*. 1–16.
  - [86] Andre Ye, Jared Moore, Rose Novick, and Amy X. Zhang. 2024. Language Models as Critical Thinking Tools: A Case Study of Philosophers. <https://doi.org/10.48550/arXiv.2404.04516> arXiv:2404.04516 [cs].
  - [87] Catherine Yeh, Gonzalo Ramos, Rachel Ng, Andy Huntington, and Richard Banks. 2024. GhostWriter: Augmenting Collaborative Human-AI Writing Experiences Through Personalization and Agency. *arXiv preprint arXiv:2402.08855* (2024).