Exploring the Potential of AI to Augment Thinking Within and In-between Meetings

Jude Rayan
jrayan@ucsd.edu
University of California, San Diego
La Jolla, California, USA

Steven P. Dow
spdow@ucsd.edu
University of California, San Diego
La Jolla, California, USA

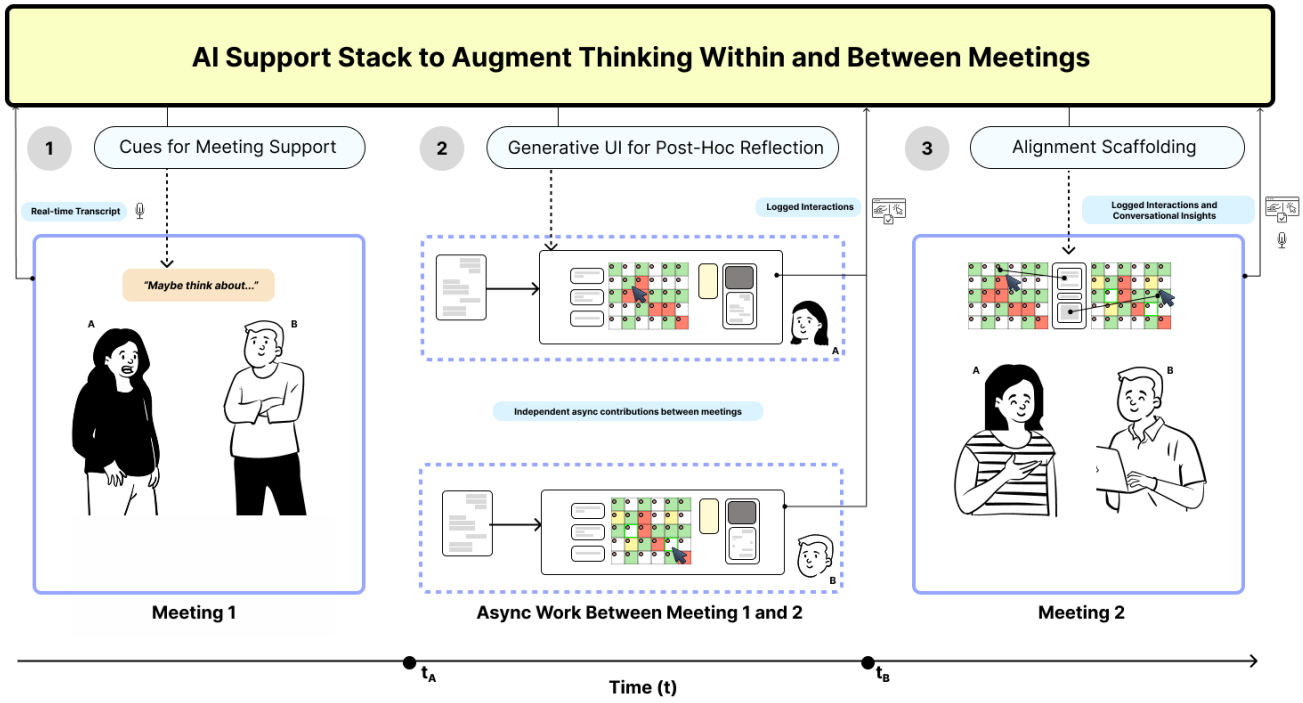

Figure 1: A snapshot of how the thinking layer supports collaboration during and between meetings. The central timeline (t) indicates the progression of time across a sequence of meetings. (1) The transcript is used to provide real-time cues to help users reach their goals. (2) After the meeting, the transcript is parsed by the layer to create Generative UI based representations that enable asynchronous, post-hoc reflection on the meeting. The system also logs user interactions and modifications within these views. t_A and t_B represent the different times at which meeting participants work on the representations. (3) This interaction data are then synthesized to surface insights and tensions that support coordination and decision-making in subsequent meetings. This cycle repeats throughout the project across multiple meetings.

Abstract

Verbal conversation is already a powerful tool for collective thinking, allowing groups to negotiate meaning, surface ideas, and build shared understanding in real time. However, meetings in general quickly unfold, lack structure, and make it difficult for groups to establish common ground or carry insights forward. Our work explores whether generative AI can be designed to amplify thinking within and between meetings. Drawing on facilitation practices, we design and study AI systems that intervene during or in-between verbal conversations. We are currently extending beyond the meeting, exploring AI-driven dynamic representations and alignment scaffolding to support the extraction, coordination, and evolution of insights across meetings. We share insights on how meeting support systems can generate shared representations that help groups reflect, coordinate, and continue thinking together.

CHI 2026 Tools for Thought Workshop, Barcelona, Spain

ACM ISBN 978-1-4503-XXXX-X/2018/06
https://doi.org/XXXXXXX.XXXXXXX

ACM Reference Format:
Jude Rayan and Steven P. Dow. 2026. Exploring the Potential of AI to Augment Thinking Within and In-between Meetings. In *Proceedings of (CHI 2026 Tools for Thought Workshop)*. ACM, New York, NY, USA, 6 pages. https://doi.org/XXXXXXX.XXXXXXX

1 INTRODUCTION

As Gen AI increasingly embeds within workflows and tools, we must consider its impact on work – not just how it shapes individual work, but how groups will interact with AI together. Prior work on computer-supported group creativity has examined collaboration across both asynchronous [18, 42] and synchronous settings [6]. To date, most of this research focuses on text-based interaction, including chatbot-based systems where AI can help users articulate, manage, and reflect on ideas [20, 41]. However, text-based systems only address a small aspect of how and where creative collaborations unfold. Real-time meetings where ideas are formed through rapid exchanges and informed by social cues are central to creative problem solving [39] yet underexplored in the literature on human-AI interaction. In our work, we focus specifically on how AI can support creativity during and in-between meetings by integrating theories from facilitation, scaffolding and coordination.

What makes real-time meetings interesting, yet difficult to support with technology? Team-mates often build on each other's half-formed thoughts through conversations, exchange information, and rapidly build on each other's ideas [24]. Being a fertile site for creative problem solving, it is loaded with subtle details around interruption, tone, timing, and gestures that impact the interpretation [1]. Non verbal behaviors such as eye gaze [12], and gestures [30] provide rich contextual cues which support trust, alignment, and collaboration. But coordinating attention, turn-taking, and idea evaluation in real time also places cognitive and social demands on participants, which can trigger well-documented social dynamics that can inhibit productivity, including groupthink [22], lurking [23, 25, 31], production blocking [16], and evaluation apprehension [15]. These dynamics can prevent groups from fully using their collective creative potential. In practice, groups often rely on experienced facilitators to help manage these challenges who use structured strategies such as managing turn-taking, inviting quieter voices to balance participation, reframe contributions, and handle conflict [33]. Yet even experienced facilitators struggle with tracking fast and overlapping speech, knowing when and how to intervene, balancing help versus disruption, and doing so without deep prior knowledge of the creative domain being discussed [8].

With Generative AI increasingly being adopted as a catalyst for creative problem solving, it opens up new possibilities for supporting creative collaboration. Yet, AI support for real-time meetings introduces unique tensions. Poorly designed interventions risk distracting participants, disrupting conversational flow, or introducing bias [32, 35, 38]. Given that creative problem solving during meetings already imposes metacognitive demands simply to participate, understanding how and when AI should act without compounding this load becomes a central design challenge [46]. We have explored this question in three parts:

(1) We studied human facilitators to understand what they attend to, how they intervene via text-based cues, and what strategies they use, in order to inform the design of AI facilitation for real-time brainstorming meetings.
(2) We examined what happens when AI intervenes in real time, asking how AI-generated facilitation cues and the modality of the meeting shape the productivity of the meeting
(3) We will be exploring how AI facilitation can extend beyond a single meeting to support creative work and coordination across multiple meetings.

1.1 What can we learn from human facilitators to guide AI facilitation for real-time meetings? (Published at ACM C&C 2024)

Prior work in HCI that uses speech to support conversations has largely focused on surfacing information about the conversation itself through systems that visualize participation patterns (e.g., Conversation Clock [10]) or highlight topics discussed (e.g., Ideation Compass [47], Conversation Clusters [11], Talk Traces [14]). While these systems treat speech primarily as something to be represented, another line of work builds on the idea of dual-purpose speech [28], where speech is both communication between people and input to computational systems. In these systems, speech is used to support task-driven conversations, often by indexing or retrieving information—for example, Serendipity Wall [21] for surfacing relevant research articles, systems that use speech to trigger actions or retrieve media (e.g., [29]), and meeting artifacts like Crosstalk [51]. Work on speech-based support for brainstorming includes systems like Inspiration Wall [3], which use spoken utterances to surface keywords. While these systems engage with conversational content, they primarily focus on idea generation rather than on facilitating the group process itself. In contrast, we were interested in treating speech not just as input for visualization or content retrieval, but as a signal to drive facilitation—supporting how people interact, participate, and coordinate during real-time verbal conversations. Therefore, we first explored what kinds of conversational cues are generated by human facilitators (role-playing as AI) during verbal ideation sessions.

Facilitating verbal group creativity is more than managing ideas. The role of a facilitator is not only to help groups generate more ideas [45], but also to manage social dynamics such as participation balance [9], turn-taking [13], managing conflict and alignment[50]. Supporting these dynamics requires more than visualizing conversational content, as it involves a facilitator's situational awareness of participant coordination during rapidly unfolding tasks—and their ability to manage attention across different parts of a shared activity [19]. Facilitators also actively reshape ideas to match evolving goals—for example, organizing ideas to surface common themes, and then reconfiguring them to help groups revisit problems from new perspectives [17]. Facilitation therefore depends not only on choosing and transforming representations, but on making them visible and actionable at the right moments. Therefore, it was important to know how different conversational cues could impact the on-going creative discussions. To take a human-first approach to understanding how AI should facilitate real-time verbal conversations through text, we turned to human facilitators.

1.1.1 Method: We built a meeting system that brought participants into a real-time verbal brainstorming session while a human wizard facilitator listened in and sent text cues through a dedicated input interface. After the conversation, The ideators reviewed each cue and reflected on its helpfulness and timeliness, while the facilitator reflected on why they sent each cue at the time and reflected on its impact on the conversation. Finally, we conducted exit interviews with the ideators and the facilitator to reflect on challenges, strategies, and opportunities for future AI-driven facilitation setups.

1.1.2 Key findings: Facilitators employed a small set of recurring strategies: (S1) dig deeper into an existing dimension, (S2) introduce a new dimension, perspective, or context, (S3) strategies for ideation, and (S4) keep the group on task and on time. Introducing new dimensions (S2) was most likely to shift the direction of the conversation. Ideators rated some strategies as more helpful and timely than others, and shared how different cues affected their thinking. Some cues drew attention to identifying additional stakeholders, others helped uncover deeper dimensions of the design space, increased the specificity of ideas, or supported critique and refinement. Across all strategies, timing of the cues was crucial. If a cue that arrived too early or too late, it was often ignored. Insufficient prior knowledge also limited ideation as some participants struggled to engage because they were unfamiliar with the topic or found it hard to get started.

Facilitators reflected that effective cueing requires close attention to conversational signals such as lulls or pauses, the specificity of ideas being discussed, and the pace of idea generation. They also described challenges around timing, domain knowledge, and phrasing. Formulating cues at the right moment was demanding, topic expertise shaped how meaningfully they could intervene, and crafting short, clear, and motivating cues under time pressure was itself difficult. Finally, we found that nearly half of all cues were missed, highlighting the need to balance timely support [43] with the attentional demands imposed by cues [5] as key considerations for successful real-time conversational facilitation. Accordingly, facilitation systems should integrate approaches that are cognizant of how interventions interact with associative memory search and attentional demands, leveraging multi-modal conversational signals to guide when to intervene without disrupting creative cognition [7].

1.2 Do real-time AI facilitation cues and the modality of the meeting have an impact on the productivity of a verbal brainstorming discussion? (Published at ACM CI 2025)

The first study surfaced attention to facilitation cues as a key factor, and it was conducted entirely in virtual settings. However, facilitation cues may function differently depending on whether interaction is virtual or in person, where participants share physical space, gestures, and eye contact that compete with or support attention. While the previous study helped us understand facilitation strategies, it did not tell us whether AI-generated cues would have any real impact on creative outcomes. If AI facilitation is to be deployed in real-world settings, we wanted to investigate whether these cues attract attention, influence participation, and ultimately impact verbal brainstorming productivity. Therefore, this depends on how facilitation cues are perceived and how they shape ongoing conversation. To make AI facilitation viable in practice, we need to know whether AI-generated cues impact verbal brainstorming productivity, and how the modality of the conversation impact attention to these cues.

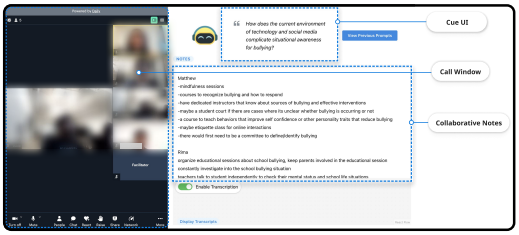

(a)

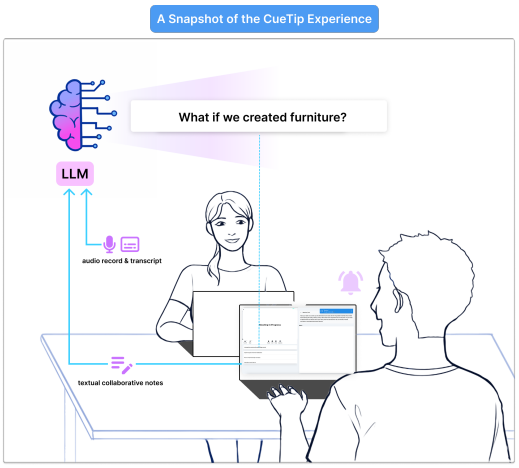

(b(i))

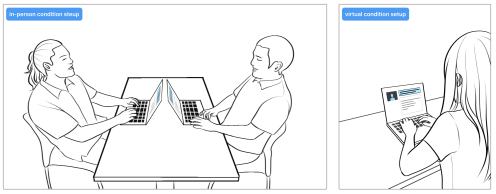

(b(ii))

Figure 2: (a) System used in Study 1, where human wizard facilitators sent text cues through a cue interface and participants received them alongside a shared real-time note-taking space. (b(i)) System used in Study 2 (CueTip), which indexed real-time transcripts from laptop microphones and generated AI-based facilitation cues. (b(ii) In the in-person condition, participants were seated together, whereas in the virtual condition, participants were located in separate rooms.

1.2.1 Method: We built CueTip, a meeting system that indexes real-time utterances from verbal conversations and generates AI cues to support ongoing discussion. Guided by findings from our

prior study, we added both a visual notification and an audio chime each time a cue was sent. We conducted a 2×2 between-subjects study in which participants were randomly assigned to one of four conditions: NoCue–Virtual, NoCue–InPerson, Cue–Virtual, and Cue–InPerson. The Cue/NoCue factor determined whether groups received facilitation cues, while the Virtual/InPerson factor determined whether participants collaborated remotely or face-to-face. Dyads in the virtual condition were placed in separate rooms, while dyads in the in-person condition sat across from each other. For experimental control, cues were generated at a fixed interval of once per minute. We measured brainstorming productivity using the number of ideas generated, number of cues attended in each modality, topical diversity of the generated ideas, and topical diversity of the conversation.

1.2.2 Key Findings: We found no significant increase in idea quantity overall, but we observed an increase in the topical diversity of both the generated ideas and the conversation. Attention to cues was positively correlated with the topical diversity of ideas, and in virtual groups, attention to cues was also positively correlated with the number of ideas generated. On average, participants in in-person conditions attended to more cues than those in virtual conditions. Qualitative findings showed that cues often helped jumpstart discussion and provide momentum during brainstorming. While cues consistently caught participants' attention, their impact ranged from helpful to distracting. Some cues supported ideation, while others disrupted flow. It is important to note that sustained conversation-driven problem solving can be cognitively fatiguing, as participants must continuously generate, interpret, and evaluate contributions in real time. When facilitation cues are not sensitive to these limitations, mental fatigue can limit participants' ability to appropriately interpret and evaluate cues, potentially lowering ideation quality over time [37]. Participants expressed preferences around how support was invoked and how notifications were delivered. They wanted cues that guided the task while dynamically adapting to their needs. Future work could isolate cue content by systematically comparing categories of prompts to understand which intervention strategies are most effective across different brainstorming scenarios. It would also be valuable to examine how these cues generalize to other meeting contexts, where conversational needs may differ substantially. More broadly, an important direction for future work is to study how both cue content and delivery can be adapted to better align with evolving meeting needs and interactional dynamics. Additionally, richer ways to evaluate the usefulness of generated ideas, beyond number of generated ideas and their topical diversity should be studied. We also need to understand how these systems work with 'in-the-wild' settings, with larger groups. Another direction is to explore other modalities for facilitation cues, such as images or sketches, rather than only text. A key limitation we observed was a tendency toward repetition in AI-generated cues, which could contribute to idea homogenization over time [2]. As a potential workaround, we explored introducing controlled noise into the generation pipeline which is yet to be evaluated. Rather than being random, these cues could be contextualized to participants' backgrounds, allowing the system to introduce variation while remaining relevant to the group.

1.3 Exploring the Potential for Generative UI from Conversations to Facilitate Group Co-ordination Across Meetings (Future Work)

After studying how AI can support real-time meetings, we are currently exploring what happens after and in-between meetings. Drawing on Star's concept of boundary objects [44], we view artifacts produced during collaboration as mediators across time, roles, and contexts. As Viller argues, collaboration is not confined to meetings or moments of interaction, but unfolds across time [49]. Members of creative groups engage in a range of artifact-level orchestration practices during and between these meetings [34]. This motivates a shift from designing AI that intervenes in the moment to designing AI that produces evolving artifacts groups can work with across meetings.

Recent work on meeting intentionality and temporal work further reinforces this need by showing that meetings often lack sustained intentionality, as goals shift, uncertainty emerges, and discussions drift over time [40, 48]. Their studies demonstrate how AI-assisted prospective, active, and passive reflection can help participants clarify purpose, assess whether discussions are "on track," and adapt plans before and during meetings while also surfacing tensions around timing, attention, and disruption. Flexible support could be beneficial in helping articulate tensions while engaging in meaningful reflective thinking, especially during collaborative decision-making [36]. Complementing this, Rintel et al.'s work on temporal work across meetings frames collaboration as spanning acts of retrospection and prospection across multiple meetings and projects, arguing that generative AI systems should be designed to support these temporal practices rather than isolated moments of interaction [48]. Together, this work highlights that meetings are embedded within longer trajectories of work, where intentions, meanings, and priorities are continuously renegotiated. Yet most collaborative systems treat representations as static outcomes such as lists of ideas, final decisions, or summaries rather than as dynamic resources that evolve with the work [4, 52]. This creates a gap between how collaboration actually unfolds and how it is represented: as goals shift, tensions emerge, and priorities change, the representations meant to support collaboration often lag behind or flatten this process.

In this framing, conversational artifacts like notes, ideas, summaries, and traces of interaction, become interfaces for coordination, interpretation, and continuity. This motivates a core question: how can systems generate interactive representations [26, 27] from shared conversational context that support both immediate coordination and longer-term sense-making? We explore pipelines that mine verbal conversation data to generate dynamic representations grounded in conversational data that reflect evolving shared mental models and coordination needs. Building on this, our future work extends facilitation beyond real-time support within meetings to examine how Generative UI can support post-hoc reflection by:

(1) accurately extracting key elements of a meeting beyond what current AI summaries provide
(2) providing scaffolding during meetings without distracting participants

(3) reducing cognitive barriers to engage in reflection and making meeting goals actionable
(4) helping coordinate alignment for subsequent meetings

2 Acknowledgements

We thank the National Science Foundation for its support of the two prior published projects referenced in this paper under Grant #2009003.

References

[1] Rima Al Tawil. Nonverbal communication in text-based, asynchronous online education. *International Review of Research in Open and Distributed Learning*, 20(1), 2019.
[2] Barrett R Anderson, Jash Hemant Shah, and Max Kreminski. Homogenization effects of large language models on human creative ideation. In *Proceedings of the 16th conference on creativity & cognition*, pages 413–425, 2024.
[3] Salvatore Andolina, Khalil Klouche, Diogo Cabral, Tuukka Ruotsalo, and Giulio Jacucci. Inspirationwall: supporting idea generation through automatic information exploration. In *Proceedings of the 2015 ACM SIGCHI Conference on Creativity and Cognition*, pages 103–106, 2015.
[4] Sumit Asthana, Sagi Hilleli, Pengcheng He, and Aaron Halfaker. Summaries, highlights, and action items: Design, implementation and evaluation of an llm-powered meeting recap system. *Proceedings of the ACM on Human-Computer Interaction*, 9(2):1–29, 2025.
[5] Brian P Bailey, Joseph A Konstan, and John V Carlis. Measuring the effects of interruptions on task performance in the user interface. In *Smc 2000 conference proceedings. 2000 ieee international conference on systems, man and cybernetics.'cybernetics evolving to systems, humans, organizations, and their complex interactions'(cat. no. 0*, volume 2, pages 757–762. IEEE, 2000.
[6] Patti Bao, Elizabeth Gerber, Darren Gergle, and David Hoffman. Momentum: getting and staying on topic during a brainstorm. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '10, page 1233–1236, New York, NY, USA, 2010. Association for Computing Machinery.
[7] Roger E Beaty and Yoed N Kenett. Associative thinking at the core of creativity. *Trends in cognitive sciences*, 27(7):671–683, 2023.
[8] Frances Bee and Roland Bee. *Facilitation skills*. CIPD Publishing, 1998.
[9] Lee Anne Bell, Diane J Goodman, and Mathew L Ouellett. Design and facilitation. In *Teaching for diversity and social justice*, pages 55–93. Routledge, 2016.
[10] Tony Bergstrom and Karrie Karahalios. Conversation clock: Visualizing audio patterns in co-located groups. In *2007 40th Annual Hawaii International Conference on System Sciences (HICSS'07)*, pages 78–78. IEEE, 2007.
[11] Tony Bergstrom and Karrie Karahalios. Conversation clusters: grouping conversation topics through human-computer dialog. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 2349–2352, 2009.
[12] Leanne S Bohannon, Andrew M Herbert, Jeff B Pelz, and Esa M Rantanen. Eye contact and video-mediated communication: A review. *Displays*, 34(2):177–185, 2013.
[13] Mike Brown. The facilitator as gatekeeper: A critical analysis of social order in facilitation sessions. *Journal of Adventure Education & Outdoor Learning*, 2(2):101–112, 2002.
[14] Senthil Chandrasegaran, Chris Bryan, Hidekazu Shidara, Tung-Yen Chuang, and Kwan-Liu Ma. Talktraces: Real-time capture and visualization of verbal content in meetings. In *Proceedings of the 2019 CHI conference on human factors in computing systems*, pages 1–14, 2019.
[15] Panayiota A Collaros and Lynn R Anderson. Effect of perceived expertness upon creativity of members of brainstorming groups. *Journal of Applied Psychology*, 53(2p1):159, 1969.
[16] Michael Diehl and Wolfgang Stroebe. Productivity loss in brainstorming groups: Toward the solution of a riddle. *Journal of personality and social psychology*, 53(3):497, 1987.
[17] Jelle van Dijk, Jirka van der Roest, Remko van der Lugt, and Kees CJ Overbeeke. Noot: a tool for sharing moments of reflection during creative meetings. In *Proceedings of the 8th ACM conference on Creativity and cognition*, pages 157–164, 2011.
[18] Haakon Faste, Nir Rachmel, Russell Essary, and Evan Sheehan. Brainstorm, chainstorm, cheatstorm, tweetstorm: new ideation strategies for distributed hci design. In *Proceedings of the sigchi conference on human factors in computing systems*, pages 1343–1352, 2013.
[19] Carl Gutwin and Saul Greenberg. A descriptive framework of workspace awareness for real-time groupware. *Computer Supported Cooperative Work (CSCW)*, 11(3):411–446, 2002.
[20] Angel Hsing-Chi Hwang and Andrea Stevenson Won. Ideabot: investigating social facilitation in human-machine team creativity. In *Proceedings of the 2021 CHI conference on human factors in computing systems*, pages 1–16, 2021.
[21] Shota Imamura, Hirotaka Hiraki, and Jun Rekimoto. Serendipity wall: A discussion support system using real-time speech recognition and large language model. In *Proceedings of the Augmented Humans International Conference 2024*, pages 237–247, 2024.
[22] Irving L Janis. Groupthink: The desperate drive for consensus at any cost. *Psychology Today*, 5(6):43–76, 1971.
[23] Norbert L Kerr and Steven E Bruun. Dispensability of member effort and group motivation losses: Free-rider effects. *Journal of Personality and social Psychology*, 44(1):78, 1983.
[24] Nicholas W Kohn, Paul B Paulus, and YunHee Choi. Building on the ideas of others: An examination of the idea combination process. *Journal of Experimental Social Psychology*, 47(3):554–561, 2011.
[25] Mestan Küçük. Lurking in online asynchronous discussion. *Procedia-Social and Behavioral Sciences*, 2(2):2260–2263, 2010.
[26] Yaniv Leviathan, Dani Valevski Matan Kalman, Danny Lumen, Eyal Segalis Eyal Molad, Shlomi Pasternak, Vishnu Natchu, Valerie Nygaard, and Srinivasan Cheenu Venkatachary James Manyika Yossi Matias. Generative ui: Llms are effective ui generators. Technical report, Technical report, Google Research, 2025. Available at generativeui. github. io.
[27] Siân Lindley, Jack Williams, Yining Cao, Haijun Xia, Elizabeth F Churchill, Abigail Sellen, Jeffrey Nichols, and David R Karger. What does generative ui mean for hci practice? 2026.
[28] Kent Lyons, Christopher Skeels, Thad Starner, Cornelis M Snoeck, Benjamin A Wong, and Daniel Ashbrook. Augmenting conversations using dual-purpose speech. In *Proceedings of the 17th annual ACM symposium on User Interface Software and Technology*, pages 237–246, 2004.
[29] Moira McGregor and John C Tang. More to meetings: challenges in using speech-based technology to support meetings. In *Proceedings of the 2017 ACM conference on computer supported cooperative work and social computing*, pages 2208–2220, 2017.
[30] Philipp Müller, Michael Xuelin Huang, and Andreas Bulling. Detecting low rapport during natural interactions in small groups from non-verbal behaviour. In *Proceedings of the 23rd International Conference on Intelligent User Interfaces*, pages 153–164, 2018.
[31] Blair Nonnecke and Jenny Preece. Why lurkers lurk. 2001.
[32] Eyal Ofek, Shamsi T Iqbal, and Karin Strauss. Reducing disruption from subtle information delivery during a conversation: mode and bandwidth investigation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 3111–3120, 2013.
[33] Anne K Offner, Thomas J Kramer, and Joel P Winter. The effects of facilitation, recording, and pauses on group brainstorming. *Small group research*, 27(2):283–298, 1996.
[34] Srishti Palani and Gonzalo Ramos. Evolving roles and workflows of creative practitioners in the age of generative ai. In *Proceedings of the 16th Conference on Creativity & Cognition*, pages 170–184, 2024.
[35] Henning Pohl, Andreea Muresan, and Kasper Hornbæk. Charting subtle interaction in the hci literature. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pages 1–15, 2019.
[36] Leon Reicherts, Zelun Tony Zhang, Elisabeth von Oswald, Yuanting Liu, Yvonne Rogers, and Mariam Hassib. Ai, help me think—but for myself: Assisting people in complex decision-making by providing different kinds of cognitive support. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, pages 1–19, 2025.
[37] Bruce A Reinig and Robert O Briggs. Putting quality first in ideation research. *Group Decision and Negotiation*, 22(5):943–973, 2013.
[38] Claudia Roda and Julie Thomas. Attention aware systems: Theories, applications, and research agenda. *Computers in Human Behavior*, 22(4):557–587, 2006.
[39] Keith Sawyer. Creativity and dialogue: The improvisational nature of conversational interaction. In *The Routledge handbook of language and creativity*, pages 78–91. Routledge, 2015.
[40] Ava Elizabeth Scott, Lev Tankelevitch, Payod Panda, Rishi Vanukuru, Xinyue Chen, and Sean Rintel. What does success look like? catalyzing meeting intentionality with ai-assisted prospective reflection. In *Proceedings of the 4th Annual Symposium on Human-Computer Interaction for Work*, CHIWORK '25, New York, NY, USA, 2025. Association for Computing Machinery.
[41] Joon Gi Shin, Janin Koch, Andrés Lucero, Peter Dalsgaard, and Wendy E. Mackay. Integrating ai in human-human collaborative ideation. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI EA '23, New York, NY, USA, 2023. Association for Computing Machinery.
[42] Pao Siangliulue, Joel Chan, Steven P Dow, and Krzysztof Z Gajos. Ideahound: improving large-scale collaborative ideation with crowd-powered real-time semantic modeling. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*, pages 609–624, 2016.
[43] Pao Siangliulue, Joel Chan, Krzysztof Z Gajos, and Steven P Dow. Providing timely examples improves the quantity and quality of generated ideas. In *Proceedings of the 2015 ACM SIGCHI Conference on Creativity and Cognition*, pages 83–92, 2015.
[44] Susan Leigh Star and James R Griesemer. Institutional ecology,translations' and boundary objects: Amateurs and professionals in berkeley's museum of

vertebrate zoology, 1907-39. *Social studies of science*, 19(3):387–420, 1989.

[45] Nawel Takouachet, Jérémy Legardeur, and Iban Lizarralde. The role of the facilitator during digital creative sessions. In *Proceedings of the 2014 Ergonomie et Informatique Avancée Conference-Design, Ergonomie et IHM: quelle articulation pour la co-conception de l'interaction*, pages 20–23, 2014.

[46] Lev Tankelevitch, Viktor Kewenig, Auste Simkute, Ava Elizabeth Scott, Advait Sarkar, Abigail Sellen, and Sean Rintel. The metacognitive demands and opportunities of generative ai. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, pages 1–24, 2024.

[47] Sander Välk, Chitipat Thabsuwan, and Céline Mougenot. The ideation compass: supporting interdisciplinary creative dialogues with real time visualization. *International Journal of Design Creativity and Innovation*, 11(2):99–116, 2023.

[48] Rishi Vanukuru, Payod Panda, Xinyue Chen, Ava Elizabeth Scott, Lev Tankelevitch, and Sean Rintel. Designing interfaces that support temporal work across meetings with generative ai. In *Proceedings of the 2025 ACM Designing Interactive Systems Conference*, pages 3600–3620, 2025.

[49] Stephen Viller. The group facilitator: a cscw perspective. In *Proceedings of the Second European Conference on Computer-Supported Cooperative Work ECSCW'91*, pages 81–95. Springer, 1991.

[50] Harry M Webne-Behrman. *The practice of facilitation: Managing group process and solving problems*. Bloomsbury Publishing USA, 1998.

[51] Haijun Xia, Tony Wang, Aditya Gunturu, Peiling Jiang, William Duan, and Xiaoshuo Yao. Crosstalk: Intelligent substrates for language-oriented interaction in video-based communication and collaboration. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, pages 1–16, 2023.

[52] Zheng Zhang, Weirui Peng, Xinyue Chen, Luke Cao, and Toby Jia-Jun Li. Ladica: a large shared display interface for generative ai cognitive assistance in co-located team collaboration. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, pages 1–22, 2025.

Received 20 February 2007; revised 12 March 2009; accepted 5 June 2009