AI and the Future of Collaborative Work: Group Ideation with an LLM in a Virtual Canvas

Jessica He
IBM Research AI
Seattle, USA
jessicahe@ibm.com

Stephanie Houde
IBM Research AI
Cambridge, USA
Stephanie.Houde@ibm.com

Gabriel Enrique Gonzalez
IBM Argentina
Necochea, Argentina
gabriel.gonzalez@ibm.com

Dario Andres Silva Moran
IBM Argentina
La Plata, Argentina
dario.silva@ibm.com

Steven I. Ross
IBM Research AI
Cambridge, USA
steven_ross@us.ibm.com

Michael Muller
IBM Research AI
Cambridge, USA
michael_muller@us.ibm.com

Justin D. Weisz
IBM Research AI
Yorktown Heights, USA
jweisz@us.ibm.com

ABSTRACT

The introduction of generative AI into multi-user applications raises novel considerations for the future of collaborative work. How might collaborative work practices change? How might we incorporate generative AI into shared tools with users' needs at the forefront? We examine these questions in the context of a remote team conducting ideation tasks – an example of collaborative work enabled by a shared digital workspace. We conducted a user study with 17 professionals experienced with virtual group ideation workshops. Our study examined their use of the *Collaborative Canvas*, a virtual canvas tool with integrated generative AI capabilities that we created as a probe. Participants saw value in using generative AI to assist with group facilitation and to augment perspectives and ideas. However, they worried about losing human perspectives and critical thinking, as well as reputational harms resulting from harmful AI outputs. Participants shared suggestions for appropriate ways to incorporate generative AI capabilities within multi-user applications and identified needs for transparency of content ownership, private digital spaces, and specialized AI capabilities. Based on participants' insights, we share implications and opportunities for the incorporation of generative AI into collaborative work in ways that place user needs at the forefront.

CCS CONCEPTS

• **Human-centered computing** → **Empirical studies in HCI**; **Empirical studies in collaborative and social computing**; • **Computing methodologies** → **Artificial intelligence**.



This work is licensed under a Creative Commons Attribution-NonCommercial International 4.0 License.

CHIWORK '24, June 25–27, 2024, Newcastle upon Tyne, United Kingdom
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1017-9/24/06
https://doi.org/10.1145/3663384.3663398

KEYWORDS

Group ideation, Brainstorming, Shared virtual canvas, Mixed initiative, Generative AI, Future of work

ACM Reference Format:
Jessica He, Stephanie Houde, Gabriel Enrique Gonzalez, Dario Andres 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 (CHIWORK '24), June 25–27, 2024, Newcastle upon Tyne, United Kingdom.* ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3663384.3663398

1 INTRODUCTION

Generative AI has the potential to automate tasks that were previously thought to be exclusive to humans, raising questions of its impact on the future of work. Rapid advancements in generative AI have already enabled its incorporation into professional work domains such as programming [95], research [1, 60], and design [41, 57, 113]. In these domains, generative AI can produce ideas, summaries, and suggestions that rival those produced by humans.

In the past, many research studies on AI-augmented work have been configured such that *one* human works with a single AI agent (e.g., [55, 94, 119, 122]). Work by Farrell et al. [26] took an expanded view in which a human works with a back-end collection of AI agents, through a single front-end interface [26]. However, this view does not reflect the reality that work is often performed by *groups* of colleagues working together [104, 116, 127]. As advancements in generative AI make their way into collaborative domains, there is a need to shift our view to realistic work contexts in which multiple humans interact with each other along with AI agents.

One indication of the rise of AI-supported group work is in the introduction of generative AI features to digital canvas tools such as Mural[1], Miro[2], and Microsoft Whiteboard[3]. These features include

[1]Mural. https://mural.co
[2]Miro. https://miro.com
[3]Microsoft Whiteboard. https://whiteboard.microsoft.com

the use of AI to generate ideas, create diagrams and presentations, summarize content, and cluster ideas into themes, mirroring ideas explored in research studies of creativity and ideation [9, 35, 54, 76]. As remote work becomes more common [6, 8, 21, 34, 51, 80], these canvas tools have become a way for teams – both remote and in-person – to have a digital collaborative experience together [19, 37]. In parallel, recent studies have also begun to explore the use of AI in collaborative tasks, such as brainstorming [30, 113, 126] and co-creation [5, 41, 105, 113, 124]. These studies have started to identify the utility and limitations of using generative AI in multi-user tasks. However, there remains a need to understand how generative AI affects group work practices in real-world settings, how to use it to enhance (rather than automate) collaboration, and how to mitigate the risks that it poses to the future of collaborative work.

To address this gap, we studied domain experts' perceptions of the unique benefits, risks, and needs that arise from the use of generative AI in a collaborative work context, and how these considerations may impact the future of their work. To explore this topic, we focused on group ideation as a case study of collaborative work that takes place in a shared digital workspace, and we used a canvas tool as an example of a shared digital workspace. We chose group ideation because it is a common task with which workers are highly familiar, and it is a task that does not require prior training. We conducted an interview study with 17 professionals who facilitate or regularly participate in group ideation workshops as part of their work. We built a canvas tool powered by a large language model (LLM) called *Collaborative Canvas* to experiment with different ways a team could interact with generative AI and leverage it in support of their group work. Our goal was to use *Collaborative Canvas* as a probe [49] to understand how domain experts, who currently make extensive use of non-AI canvas tools for collaboration, might engage with and view novel forms of AI support in their future work.

Our paper makes the following contributions to the HCI community:

- Implications for the future of AI-supported group ideation based on insights from domain experts, including the values of AI in augmenting perspectives and assisting with facilitation, along with the liabilities it poses to human ideas, collaborative processes, and workplace impressions. Building on these findings, we illustrate an example of how humans and AI can dynamically allocate initiative in a collaborative context to draw on the strengths of AI support while mitigating its risks.
- User needs and opportunities for incorporating generative AI capabilities into future collaborative workspaces in human-centered ways, including transparency of content ownership, private digital spaces, and specialized AI capabilities.

2 RELATED WORK

Our work empirically examines users' experiences with a multi-person, AI-infused virtual canvas. We first review prior work with graphical canvases and related media. We then consider what has been learned about brainstorming with an AI partner, primarily in dyadic, one-human-one-AI configurations. Finally, we reflect that new collaborative environments consisting of multiple humans and AIs leads to new challenges in mixed initiative interfaces.

2.1 Human Collaboration through Graphical Media

Researchers and practitioners have studied human-human collaboration through graphical tools [59] such as mindmaps [11, 12], whiteboards [72, 106], affinity diagrams [47, 68], and a series of sociotechnical templates in Total Quality Management [43, 109]. Following successes with physical media, researchers began to experiment with computer-mediated graphical tools for both same-place and virtual-place collaborations [11, 61, 69, 100]. In some cases, graphical collaboration tools took the form of *systems for argumentation*, such as gIBIS [16], Design Space Analysis [70], and Beyond Discussions [24], as well as more recent experiments such as nodenogg [92]. Others focused on sticky notes as both physical and virtual collaboration tools – Christensen et al. [15], Ho and Tomitsch [46], and Sibbet [103] provide collections and surveys of studies and toolkits of these systems. For example, earlier practices such as contextual inquiry and design thinking showed how sticky notes could serve as a backbone infrastructure in face-to-face collaborations for design, analysis, and the social construction of agreed solutions [2, 47, 121].

Lutters [69] noted two decades ago that the contents of hand-written, physical *micronotes* were unavailable to organizational memory systems (see also Christensen and Friis-Olivarius [14] on "cognitive units in the form of small structured ideas"), indicating a need to selectively store and share micronotes – possibly through virtual sticky notes. Responding to needs of geographically-distributed institutions, Gumienny et al. [36] described the benefits of online sticky notes for both ideation and feedback, with particular value to *remote* members of teams with a preponderance of members in one location. Jensen et al. [50] compared the use of physical and digital sticky notes, and concluded that digital media could support collaborative ideation. In a different domain, Briggs et al. [5] reached a similar conclusion in a more complex graphical+gestural environment, as did Cvetkovic et al. [18]. Within these types of digital workspaces that use virtual sticky notes, work by Takouachet et al. [110], Hilliges et al. [45], and Christensen and Abildgaard [13] has explored users' practices and needs for working effectively in these tools. Chen et al. [12] studied how graphical materials could be used for two-person structured collaborations, and subsequently Chen and Krishnamurthy [11] added AI support for one-human-one-AI interactions based on ConceptNet [107].

It is now commonplace for teams to use graphical collaboration tools with virtual sticky notes, such as Miro, Mural, and Microsoft Whiteboard. Developers of those tools are now adding generative AI capabilities [44, 74, 93], prompting us to examine the role of these capabilities and their impact on the collaborative dynamic, such as in group ideation.

2.2 Brainstorming with Generative AI Applications

Group ideation, also known as brainstorming, is the process by which a group of people come together to generate new ideas.

Osborn [82] performed early empirical analyses of human-human brainstorming and highlighted two major phases (see also [29]):

- **Divergent thinking**, in which team members create many heterogeneous ideas, and
- **Convergent thinking**, in which members critique, combine, and select a subset of ideas for recommendation or for further development.

Other researchers have further developed these ideas, materials, and practices, creating variations on Osborn's work (e.g., [39, 40, 85, 87, 88]). While laboratory-style studies of unrelated individuals have shown a collaboration penalty for brainstorming in groups [7, 81, 90, 108], studies of teams and mutually-committed groups have shown increased productivity and conceptual richness through exposure to each other's ideas (e.g. [22, 27, 53, 86, 96]), especially when online tools allow for parallel activities to occur during brainstorming [29, 66, 75, 84]. Practices in design thinking, for example, have adopted physical sticky notes and their virtual analogues as core sociotechnical resources [20, 41, 89, 98].

Early explorations have simulated or added AI functionality to these ideation tasks and studied their impacts on users. Using an experimental method with a single, consistent script of activities, researchers have found inconsistent impacts on human responses when they labeled the scripted activities as performed by a human or an AI agent [31, 120]. Going beyond simulation and instructions, Muller and colleagues showed plausible co-creativity when one human ideated with an AI partner [76, 77, 79]. Yu-Han and Chun-Ching [126] found enhanced human performance when working with a LLM, and Bouschery et al. [3] similarly found superior performance when one human brainstormed with one AI agent. Lavrič and Škraba [63] described somewhat contrasting outcomes based on human evaluation of ideas generated by GPT-3.5-turbo[4] at different temperature settings[5]. In design, Tholander and Jonsson [113] found that LLMs were useful for inspiring designers with new ideas, but their lack of contextual understanding and depth of ideas spoke to the irreplacability of humans in creative work. Maier et al. [71] examined a slightly different use case – the *facilitation* of an ideation process – and described humans' overall preferences for a human facilitator over an AI facilitator.

The use of AI-generated ideas within a brainstorming session, or even the simple exposure to AI-generated ideas, raises new issues. Rezwana and Maher [94] identified ethical dilemmas in the ownership, accountability, and leadership of co-creative work. Authors of creative works, including writers and designers, have shown preferences for AI to act in more limited, supportive roles [32, 58]. Lehmann et al. [64] found, in the context of shorter texts (such as those produced in brainstorming), AI-generated materials led to reduced activity by human ideators, but a strong sense of human ownership of the resulting ideas. Despite claims of the emancipatory role of AI assistance in ideation tasks [115], Yin et al. [125] identified concerns around the reduction of human-generated texts, suggesting that people may be discouraged if there is too much activity by the AI agent. Laine concurred, stating that, "AI... should remain in the background, allowing the human to make decisions and retain control of the creative process" [62, p. i]. One way to maintain human control may be through relatively short conversational turns [76, 95] or actions [67] by the AI to allow the human to more frequently re-assert guidance and control.

While recent work examines the use of AI in collaborative tasks, it focuses primarily on considerations and challenges arising from interactions between one human and an (often external) AI agent. In this paper, we extend this work by examining unique implications arising from a *group* of users brainstorming together in a shared workspace embedded with generative AI capabilities.

2.3 Mixed Initiative Interfaces

Previous work implicitly engaged with issues of initiative (i.e., the relative allocation and agency of human and/vs. AI partners in brainstorming activities). The inclusion of an interactive AI agent in a group setting requires new attention to concepts of initiative and mixed initiative interfaces [48]. Early research considered the allocation of function to human initiative vs. computer initiative along a uni-dimensional trade-off dimension [28, 83, 99]. Subsequently, Shneiderman proposed a two-dimensional framework in which human initiative and AI initiative ("autonomation") could be designed independently of one another, allowing classification into one of four quadrants: high-human high-AI; high-human low-AI; low-human high-AI; low-human low-AI [102]. Using a similar categorical approach, McComb et al. [73] proposed their own two-dimensional framework of quadrants based on AI proactivity: *mode* (high/proactivity vs. low/reactivity) and *focus* (process vs. product).

These categorical frameworks were good descriptions of relatively static configurations of the user interface, in which (a) usage patterns were primarily invariant and (b) users performed a single action, or a single family of similar actions, with the product - e.g., in "human in the loop" applications, in which "Humans can... directly accomplish tasks that are hard for computers" while serving as components "in the pipeline with the help of machine-based approaches" [123, ms p.1]. Muller and Weisz [78] revised these frameworks for complex organizational workflow applications in which usage patterns could vary from one workflow step to the next, and in which users could choose their own sequences of activities according to task demands and human skills. The design of the workflows allowed users to choose their position along the two dimensions of human-initiative and AI-initiative, thereby allowing them to become in-the-moment co-designers of their work within the limits of the workflow options. In this paper, we consider the case of groups of humans interacting with an AI agent, and we anticipate the emergence of additional, in-the-moment dynamic shifts in human-initiative and AI-initiative.

3 COLLABORATIVE CANVAS

To conduct our research, we built the *Collaborative Canvas*, a virtual canvas that allows multiple users to brainstorm with generative AI capabilities. We were motivated to build our own prototype, rather than using an existing commercial product, due our need to have full control over its functionality, maintain flexibility in implementing new forms of UX interactions, and connect it to our institution's internal generative AI platform. Our prototype, shown in Figure 1, offers basic functionality commonly found in shared canvas tools:

[4]https://openai.com/blog/gpt-3-5-turbo-fine-tuning-and-api-updates

[5]Temperature refers to a LLM parameter that controls the variability of generated content. Higher temperatures often result in more variable, and potentially unusual, outcomes.

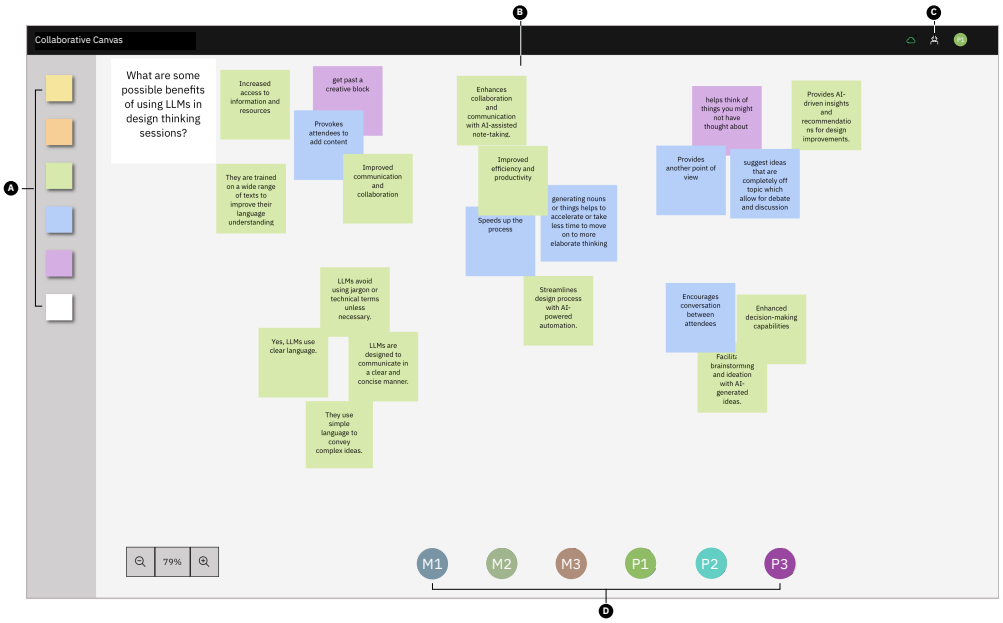


Figure 1: *Collaborative Canvas* user interface. The shared canvas provides a space for multiple users to add, edit, and remove virtual sticky notes. (A) Users can add new sticky notes in various colors from a sidebar. (B) Sticky notes can be placed and moved around the canvas. (C) Clicking on the robot icon reveals a private panel for generating sticky notes with the AI (shown in Figure 3). (D) User profile icons represent the users who are currently present in the canvas.

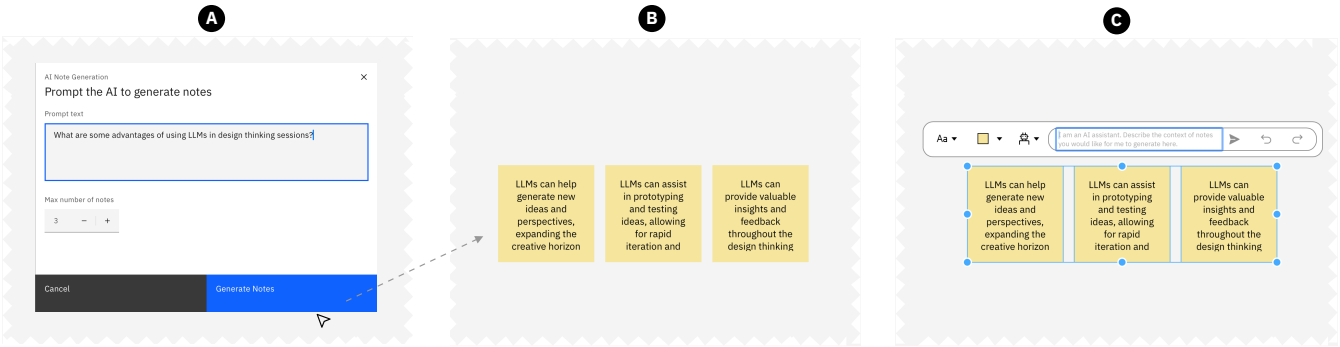


Figure 2: Generative AI capabilities within the shared canvas. Users can generate sticky notes directly in the shared canvas view. (A) Right-clicking on the canvas and choosing "Generate" opens a modal dialog in which the user can specify a prompt and request a specific number of sticky notes. (B) After clicking "Generate Notes," new sticky notes are created in the shared canvas. (C) When selecting one or more sticky notes, a pop-up toolbar allows users to re-generate the selected sticky notes with a new prompt. In this way, users can iterate on human or AI-generated sticky notes (e.g. "provide a condensed summary" or "come up with new ideas").

users can manually create virtual sticky notes, type content onto them, edit them, and delete them. We intentionally limited content modalities to only those required for basic collaboration (i.e. text and sticky notes) to minimize barriers of entry to the tool and focus sessions on the integration of generative AI. Users could prompt an LLM to generate sticky notes in two ways: through a modal dialog that generates and places new sticky notes directly on the shared canvas (Figure 2), and through a private side panel where new notes are generated and staged for the user (Figure 3). In both views, users can iterate on AI-generated notes by editing them manually or by sending a follow-up prompt to the LLM (Figures 2C and 3B). We implemented these different means for creating and refining notes so that we could study if, when, and how users chose to invoke AI assistance in a collaborative context.

We implemented *Collaborative Canvas* using the Svelte[6] framework for the front-end UI and Python with FastAPI[7] for the backend.

[6]Svelte. https://svelte.dev
[7]FastAPI. https://fastapi.tiangolo.com

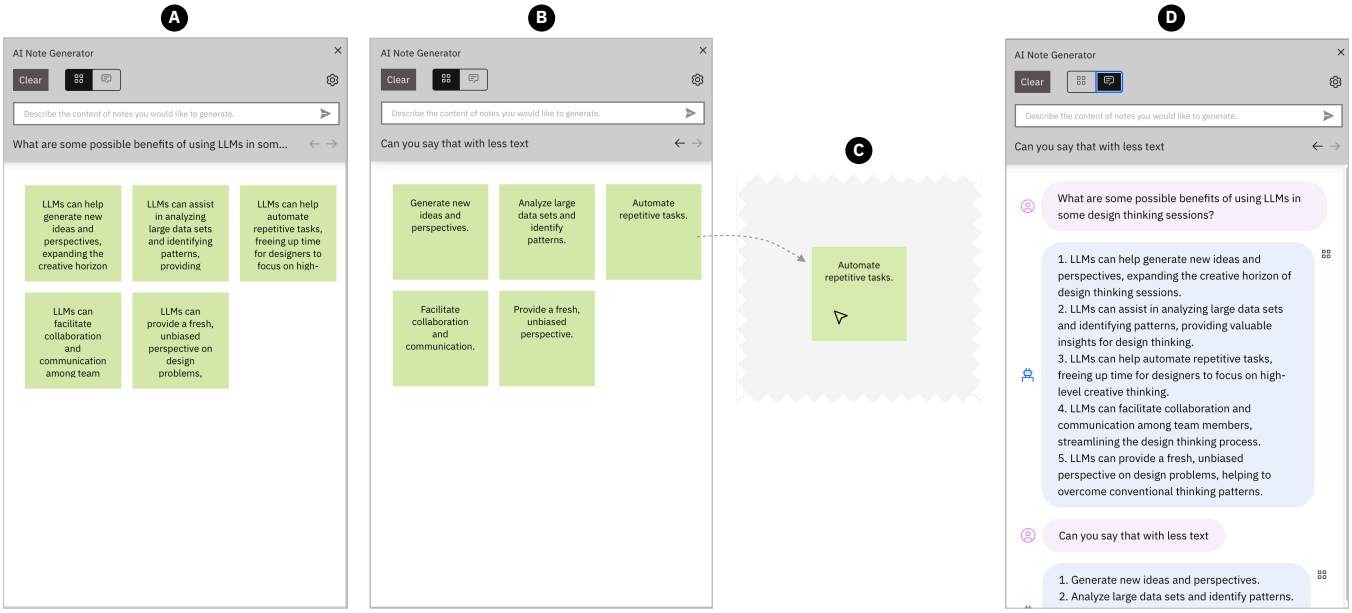


Figure 3: Private panel. As an alternative to generating in the shared canvas, we also developed a private panel for generating sticky notes. This panel allows users to review the content of generated notes before sharing them with the group. (A) The panel provides a private space to prompt the LLM and generate new sticky notes. (B) Context is preserved between subsequent prompts, so a user can follow up on prior outputs (e.g. "Can you say that with less text"). (C) Users can choose sticky notes to share with the group by dragging them from the panel to the main canvas. (D) Users can view a history of their prompts and outputs in a sequential view, where they can also click the multi-square icon to restore previously-generated notes.

We used an internally-hosted Llama 2 (70B) model [114] to power the generative AI capabilities, as we found it to perform best in conversational use cases out of the models available in our organization. Our application's backend was responsible for communicating with the internal platform that hosted the model. We developed a prompt that set the context for a brainstorming task and possible user actions, including generating sticky notes and progressively refining generated responses. This prompt also instructed the model to format its responses in a way that could be parsed into separable chunks by the front end to create multiple sticky notes from a single LLM response.

4 USER STUDY

Using *Collaborative Canvas* as a probe, we conducted an interview study with groups of professionals who facilitate or participate in group ideation workshops as part of their work. We felt that their expertise with collaborative work in shared digital workspaces would allow them to reflect on the potential impact of incorporating novel generative AI capabilities into their work. Our aim was to explore the following research questions:

- **RQ1.** How might the incorporation of generative AI into shared digital workspaces impact collaborative work practices such as group ideation?
- **RQ2.** What unique user needs arise from the incorporation of generative AI into collaborative work?
- **RQ3.** Based on these user needs, what implications are important to consider in the incorporation of generative AI capabilities into future collaborative workspaces?

4.1 Participants

We conducted seven study sessions with groups of 2-3 participants, for a total of 17 participants. The majority of participants (94%) resided in the United States; 6% resided in Mexico. Participants reported gender identities including men (47%), women (35%), and non-binary (6%); 12% preferred not to say. Participants were full-time design and software professionals within our organization, a large international technology company. All participants actively used shared canvas tools for group ideation workshops as part of their work, either as a facilitator, participant, or both. Participants had no prior experience with canvas tools that incorporated generative AI capabilities. Their familiarity with LLMs was mixed – none of them used or developed LLM technologies as part of their daily work, but some had prior exposure to generative AI applications. Their job roles and experience with canvas tools are shown in Table 1. Participants are identified as *Ps.p*, where *s* is their session number and *p* is their participant number in that session.

We recruited participants using a call for participation posted in internal Slack[8] channels for employees interested in group ideation practices. Interview groups were formed based on schedule availability. Each participant was compensated the equivalent of $25 USD for their participation.

[8]Slack. https://slack.com.

Table 1: Participant roles and use of shared digital workspaces in their current work

P#	Job Title	Use of shared digital canvases
P1.1	Visual Designer	Project stakeholder collaboration and alignment
P1.2	Design Manager	Team collaboration and alignment
P1.3	Development Manager	Alignment between designers and technologists
P2.1	UX Designer	Remote design thinking sessions with team
P2.2	Design Manager	Training and practice for design team
P3.1	UX Designer	Group analysis of business use cases
P3.2	UX Designer	Exploratory stage design with team
P3.3	UX Designer	Stakeholder alignment for business transformation
P4.1	Application Developer	Group and personal project planning
P4.2	UX Researcher	Design thinking workshops and ideation
P4.3	UX Researcher	Design thinking workshops and research analysis
P5.1	UX Researcher	Design thinkng workshops and personal project management
P5.2	UX Researcher	Design thinking workshops and project management
P6.1	Service Designer	Project stakeholder alignment
P6.2	Learning Facilitator	Teaching and presentations
P7.1	Visual Design Researcher	Research synthesis and stakeholder alignment
P7.2	UX Designer	Design thinking workshops and feedback sessions

4.2 Method

Study sessions lasted one hour and took place remotely via video conferencing. One researcher guided participants through each session while two others took notes and asked clarifying questions. Sessions consisted of four parts. First, we asked each participant to spend a few minutes describing their job role and how they use virtual canvases for group ideation in their work. Next, we demonstrated the features of the *Collaborative Canvas* and invited participants to join a shared canvas, where they spent approximately 10 minutes trying the tool for themselves and asking questions. Once everyone was familiar with the *Collaborative Canvas*, we asked participants to use it for an ideation task similar to those in their work. Finally, we held a 30-minute discussion for participants to reflect on the generative AI capabilities.

In the ideation task, participants were asked to brainstorm on this question: *"What are some possible benefits of using LLMs in design thinking sessions?"* They were given three minutes to note down as many ideas as possible in the canvas, either by manually writing their own ideas, using the integrated AI tools to generate ideas, or both, according to their individual preference. When time was up, we asked participants to cluster the ideas by theme. The goal of this task was to give participants a sense of how generative AI might integrate into a real-world group ideation workshop.

We then engaged them in a semi-structured discussion of their experience with the AI capabilities. We asked each participant to describe their approach to the task, including any integrated AI tools they used, the order in which they used them, and why they chose their approach. We also asked them to discuss any additional ways they wished that the AI could support group ideation and their thoughts on how the AI and its content should be represented. Lastly, we asked participants to speculate on the impacts – both positive and negative – of integrating generative AI into existing canvas tools and their own collaborative work practices.

5 RESULTS

Data collected in each study session included 1 hour of video recording captured with participants' consent, audio transcriptions of those recordings, contemporaneous notes taken by researchers in attendance, and the content of sticky notes created by participants during the ideation task. Two researchers collaboratively grouped and labeled the textual content of transcriptions, observer notes, and sticky notes. Using reflexive thematic analysis [4], we constructed themes and subthemes to describe participants' desired workflows with the AI, the benefits and risks they perceived, and their needs for working safely and effectively with the AI in a group context.

We begin by describing participants' experiences using the *Collaborative Canvas* for their ideation task in Section 5.1, which informed their subsequent expectations and reflections. To address **RQ1**, we analyzed participants' usage of the canvas and responses to interview questions to understand their perceptions of the impacts – both beneficial and harmful – of introducing generative AI to group ideation (Sections 5.2 and 5.3). We then identify novel user needs that emerged from their use of generative AI capabilities in a shared tool to inform **RQ2** (Section 5.4). We discuss implications of these results in the Discussion to address **RQ3** (Section 6).

5.1 Group-AI ideation practices

Participants brainstormed in a variety of ways during the task portion of the study. We present two exemplars that highlight different approaches:

(1) P2.2 *"tried to come up with as many [ideas] on [their] own before including the AI."* They first wrote down three ideas on the shared canvas (Figure 1B). Next, they accessed the generation dialog on the shared canvas (Figure 2A) and entered the task question as the prompt. After reading the AI's outputs on the canvas, P2.2 was inspired with a fourth idea, which they wrote onto a new sticky note.

(2) P4.2 began with AI support in the private panel (Figure 3A), where they entered the task question as the prompt. They dragged one AI-generated sticky note onto the canvas as-is, then selected it and opened the pop-up toolbar (Figure 2C) to ask for more specific examples of that note. Finally, they manually wrote notes in the canvas (Figure 1B) using inspiration from AI outputs and other participants' ideas.

Participants ideated in parallel, meaning multiple people wrote notes or used AI capabilities simultaneously. Many participants (58.8%) began by adding their own ideas to the shared canvas. When asked why they chose this starting point, P2.2 said they *"didn't want their thoughts to be clouded [by the AI]."* In contrast, participants who started with AI used it to learn more about the task domain or to explore its capabilities. For example, P5.2, who had limited experience with generative technologies, entered the prompt, *"who uses LLMs...why LLMs are bad...what is unethical AI?"*

When using AI for the ideation task, several participants began by entering the task question as a prompt to the AI; a few tried modifying the wording of the prompt after receiving outputs that were too generic or similar to others' outputs, with mixed success. For example, P3.3 adjusted their prompt from *"benefits of using LLMs in brainstorming"* to *"benefits of using AI."* After reviewing the new outputs, they said, *"we're onto something, but it's not quite what I'm looking for."* To generate notes, most participants preferred using the private panel so that they could select which ones to share; P3.3 said they *"don't want everything added before [they] know what it is."* Those who preferred to generate notes directly onto the canvas felt that it was simpler to use and more akin to how they currently work in virtual canvases without AI support. After receiving outputs from the AI, some participants included them on the shared canvas without modification, while others wrote their own notes after being inspired by AI-generated ideas. Participants then converged and manually clustered both human-written and AI-generated notes.

Finally, in reflecting on their experience using the *Collaborative Canvas*, some participants remarked on the novelty of incorporating AI into collaborative work. P7.2 said, *"it'd be a challenge to my status quo...I feel like my behavior with [AI] is very alone...it's just always been something I've done as an individual person, non-collaboratively, so it's challenging my perception of how I use these tools."*

5.2 Perceived benefits

Participants' comments and our observations of their use of *Collaborative Canvas* revealed several ways in which generative AI could add value to group ideation. We organize these values into three themes: providing missing expertise, overcoming lulls in ideation, and assisting with facilitation.

5.2.1 Providing missing expertise. Three participants who were less familiar with LLMs began the ideation task by asking the AI to explain different aspects of the technology, including its benefits, uses, and issues. When asked why they chose this approach, P5.1 said, *"I kind of...know enough to ask the right questions, but don't know...the answers...off the top of my head."* In line with the notion of using AI to learn about an unfamiliar domain, several participants thought the AI could provide missing perspectives and knowledge; P7.2 referred to AI as *"an expert friend to bounce your ideas off of"* and said that as a designer, they would *"lean on generative AI to plug technical gaps."* Similarly, P5.1 thought the AI could help identify use cases in an unfamiliar problem space. P2.3 thought the AI could be particularly useful in a smaller workshop to *"augment and generate more ideas,"* and in the ideation task, P3.3 wrote that the AI *"can provide diverse opinions/ideas if the group is less diverse."* P3.3 also speculated that *"AI can help identify patterns and connections that humans may miss"* during convergent thinking.

5.2.2 Overcoming lulls in ideation. Participants noted that lulls in ideation are a common pain point experienced by facilitators, either at the start or in the middle of a workshop. After experimenting with the AI's ability to generate ideas in the shared canvas and in the private panel, participants envisioned that it could mitigate this pain point by *"[getting] the ball rolling"* (P2.1) or *"moving a session along"* (P1.1). For example, P2.3 felt that the AI would be particularly helpful for kicking off discussions on sensitive topics due to its impersonal nature. P7.1 envisioned *"not using generation all the way through, but as a starting point"* when stakeholders struggle to come up with ideas. In their own ideation task during the study, P4.2 and P5.1 both used AI-generated content as inspiration for their own ideas, indicating the practicality of turning to the AI when stuck. In discussing the impact of using AI-generated ideas in these scenarios, P1.2 felt that fewer lulls would *"improve the mood and how people feel about the outcomes"* and prevent *"down moments."*

5.2.3 Assisting with facilitation. When asked about additional AI capabilities that would be useful, participants speculated that the AI could support facilitation work. One opportunity for support was help with keeping a session on schedule; for example, P1.3 envisioned that the AI could announce next steps to the group. Another opportunity was in AI-supported synthesis of content on the canvas during convergent thinking. Six participants wanted the AI to help cluster or summarize sticky notes: P1.3 wanted the AI to act as a *"smart collector,"* and P4.3 envisioned invoking this support by highlighting a few sticky notes and asking the AI, *"what's the theme within this cluster?"* P1.1 felt that AI-supported synthesis would be particularly helpful in workshops with a lot of participants, when there is a wealth of ideas and not enough time to review each one. Beyond the content on the canvas, P7.2 was also interested in the AI *"watching for rabbit holes...like an AI facilitator who's telling me, oh, you guys are going to spend too long in this session."* Finally, participants also identified pre- and post-workshop tasks in which generative AI capabilities could be valuable. For example, P7.1 wanted AI support with formulating questions for technical stakeholders when planning a session, and P7.2 wanted the AI to draft summaries and reports after a session.

5.3 Concerns and perceived liabilities

Participants discussed ways in which generative AI could harm collaborative work and instances when AI use would be particularly worrisome. From these discussions, we describe five categories of concerns and liabilities of introducing generative AI to group ideation: loss of empathy and critical thinking, homogeneity of ideas, quality issues and reputational harms, de-skilling, and the facade of human replacement.

5.3.1 Loss of empathy and critical thinking. Participants in four sessions expressed significant concerns that generative AI would undermine the unique emotional and critical thinking abilities that humans bring to collaborative work. They felt that there are unique values derived from the *process* of problem-solving, reflection, and human-to-human collaboration that AI cannot replace. As P5.2 explained, *"we have something that AI doesn't have, which is empathy, and we want to make sure that we don't water that down or take away from that."* In line with this idea, P6.1 felt that AI usage would be particularly unsuitable in ideation centered around participants sharing their feelings on a topic, such as team retrospectives. They worried that they would be *"robbed of that reflective practice, which...is really important in idea generation."* P6.2 expressed a similar sentiment on detracting from human problem-solving skills: *"I feel like if...I'm putting things into an AI generator and saying, hey, summarize this, I feel like I'm not critically thinking. I feel like I'm missing out on my chance to actually work through this problem."*

Although AI may produce outputs that rival human levels of performance, the value of group ideation often comes from the process rather than the outcomes. As P6.1 describes, *"It's often not what is generated on the board that matters, it is the collective shared experience and understanding."* P6.1 gave an example of using AI to generate a value statement: although the end result when using AI may be similar to what their team would have written without it, they would have lost out on the shared refinement journey that is critical to reaching alignment and internalizing the statement. P5.1 experienced this effect during their ideation task, when they found it more challenging to cluster AI-generated ideas due to their disconnect from the context.

5.3.2 Homogeneity of ideas. Another prevalent concern was that AI-assisted generation would produce a homogeneous set of ideas that lacked the creativity, authenticity, and diversity that a group of people could provide. Following the tutorial of *Collaborative Canvas*, P6.1 said, *"seeing the walk-through...freaks me out quite a bit because when I'm facilitating, I want ideas from people in the room...I want to get their specific perspective, sometimes in literally their words because there's metadata in how they choose to phrase things."* Echoing these concerns, P3.2 noted that idea generation assistance should be limited to facilitators, as they want workshop participants *"to come up with authentic and organic ideas from their experiences."* These comments indicate that unrestricted access to generative capabilities may be unsuitable during divergent ideation, when value is derived from each group member's thoughts and experiences. Similarly, although participants expressed a desire for synthesis support during convergent thinking (Section 5.2.3), P5.1 gave the caveat that the AI may *"flatten some of the different ideas"* by overlooking outliers that could have sparked new insights.

Concerns about homogeneity were also driven by participants' observations that the AI generated the same output for different people, sometimes even when the prompt was modified. This observation induced worries that using AI might skew a group's ideas in a certain direction by hindering creative thinking and novel ideation. P7.2 noted that the AI outputs appeared *"generic and kind of high level,"* and they worried that writing and refining LLM prompts would detract from time that workshop participants could have spent brainstorming with their unique perspectives and expertise. Reflecting on these issues, P7.2 said, *"AI is more of a risk in the creation space,"* and P6.2 said, *"we don't need to introduce [AI] at every step, only when we're really stuck."*

5.3.3 Quality issues and reputational harms. Participants raised concerns of receiving hateful, abusive, or profane (HAP) outputs based on their prior experiences with generative AI. This risk is likely exacerbated in a group setting in which other users have a shared view of any undesirable outputs. P1.3 noted that at times, AI outputs may contain embarrassing content that is unsuitable for sharing. P3.1 was especially apprehensive about hallucinations in high-stakes contexts where clients are involved, as they *"don't want to put the fate of projects on a tool that could get something wrong."* As a facilitator, P7.1 was concerned that the AI may generate *"red herrings"*: irrelevant notes that may distract stakeholders and throw their workshop off track.

5.3.4 De-skilling. Reflecting on longer term consequences, some participants worried that they would lose valuable skills involved in ideation, particularly divergent thinking. P5.1 said, *"I think one of the things that might happen is I just get lazy and ask the AI to generate...and even if you have all these adverse experiences, it spits out the same stuff."* Similarly, P6.2 said they would feel *"robbed of the experience of creating...brand new ideas if [they] just leaned on AI for everything."* To address these risks, P4.3 proposed that workshop participants should always write down their own ideas first and only use AI afterwards to cross-reference or identify new ideas, echoing sentiments of limiting AI usage expressed by participants in Sections 5.2.2, 5.3.2 and 5.3.3.

5.3.5 Facade of human replacement. Taken together, the previous categories of AI liabilities highlight the risks of believing humans are replaceable in group ideation work. As participants pointed out, the fluency of AI-generated outputs may give the impression of high quality content at first glance, but the reality is that they lack the creativity, nuance, and lived experiences that a human participant could provide. In the context of their design work, P6.1 described this risk as a *"design thinking theater,"* in which a canvas laden with AI-generated sticky notes provides the semblance of a group ideation workshop, but the value and meaning derived from human-to-human ideation are missing altogether.

5.4 User needs for shared generative AI capabilities

Participants shared ideas of design interventions that would make generative AI capabilities safer and more useful to collaborative work. While they described numerous needs ranging from minor usability improvements to novel AI support, we highlight three categories that were unique to the context of collaborative work: attribution of AI-generated content, boundaries between shared and private notes, and domain specificity and contextual knowledge.

5.4.1 Attribution of AI-generated content. Following the ideation task, a majority of participants (70.6%) commented that AI-generated notes should be differentiated from human-written notes in the shared canvas view. P1.1 thought it would be *"creepy"* to represent the AI as another human, and P6.2 thought it wouldn't be *"fair"* or *"honest"* to give others the impression that they had written

an AI-generated note. In the context of a client workshop, P1.3 thought AI attribution would be particularly important so that they wouldn't try to probe clients for reasoning behind ideas that turned out to be AI-generated.

Participants offered different ideas for how AI-generated ideas should be represented distinctly, including tagging AI-generated notes with an icon (P1.2, P2.2), outlining them (P2.3), and right-clicking them to reveal attribution information (P4.3, P7.1, P7.2). Due to the possibility of iterating on generated ideas, visual treatments may also be needed to credit collaborative human-AI efforts. For example, P2.2 proposed that if a user modified an AI-generated note, that note could be marked with an AI icon and a plus symbol. In cases where an AI-generated note is modified, P1.2, P4.3, and P7.1 also wanted a way to trace back to the original output and the prompt used to generate it. Even for AI-generated notes that had not been edited, P6.2 and P7.2 emphasized the importance of framing AI attribution as *"helped by AI"* rather than full AI ownership over the idea.

Participants' ideas for attribution raise questions of ownership in human-AI co-created content. Participants mostly felt that ownership was shared between the human and the AI. As P6.1 explained, *"The way [the user] engineered the prompt is going to influence the output. The question they ask is biased through their lens."* However, there is nuance in the extent of ownership between each party. P6.1 also discussed factoring in the significance of a user's contribution to an AI-generated note: *"if they've just fixed a typo or if they've changed the meaning of the initial note."* P3.3 felt that as prompts become more prescriptive and complex, the user is increasingly responsible for the resulting outputs, and AI attribution may become less relevant as control and responsibility over outputs shift to users.

5.4.2 Boundary between shared and private notes. Participants who generated notes directly onto the shared canvas (P5.1, P6.1, P6.2, P7.1) felt that this modality was simpler to use, as generating notes in the panel, then dragging them onto the canvas *"creates an extra step"* (P6.1). However, many participants (64.7%) appreciated the option to generate notes in a private panel, as it served as a means for them to control which AI outputs were shared with others.

The *friction* created by this boundary was desirable in certain use cases. For example, this friction may be useful to address concerns of problematic or harmful outputs (Section 5.3.3). It may also be required in high-stakes situations or to avoid reputataional harms if harmful outputs were shown to a customer or client. In reflecting on their preference for the private panel, P3.3 said they *"can be more selective about how [they're] contributing,"* and P7.2 liked *"the ability to experiment without people looking over your shoulder."*

5.4.3 Domain specificity and contextual knowledge. When asked about additional desired capabilities, six participants (35.3%) said generative AI could provide greater value to group ideation if it possessed domain-specific knowledge, context on the team's work, or awareness of content on the canvas. P4.2, P5.2, and P6.2 thought it would be useful if the AI could support tasks specific to their design thinking work, such as drafting empathy maps and problem statements. P4.2 and P4.3 wanted the AI to have data on their organization's prior research findings so that it could help stakeholders access insights more easily and generate ideas grounded in their prior work. P2.3 proposed, *"it might be interesting if [the AI] could learn from...the human ideas that were added...if it would read those and then it was like, oh, they're looking for this and come up with three others."*

6 DISCUSSION

Participants' experiences with *Collaborative Canvas* and their discussions of its AI capabilities provided insight into their desired ways of working with generative AI in a group setting. Their perceptions of the usefulness of AI in selectively augmenting ideas and helping with facilitation, as well as their concerns about AI taking away from the richness of human ideation and damaging workplace impressions informed **RQ1**. We addressed **RQ2** by describing participants' needs for mitigating those risks and working safely with the AI, including human-AI attribution, boundaries between shared and private content, and context specificity. To answer **RQ3**, we discuss people's desired roles for generative AI in future group ideation workflows, along with emerging changes to consider in safely and effectively incorporating such technologies into collaborative workspaces.

6.1 Desired role of AI in group ideation

One way to understand participants' desired roles for generative AI is through the lens of initiative: who has initiative, and for which tasks? In Figure 4, we show a dynamic representation of participants' desired AI support at different stages of a group ideation workflow, plotted within a two-dimensional space representing varied degrees of human and AI initiative. This framework is based on an extension to Shneiderman [101]'s framework of human-centered AI proposed by Muller and Weisz [78] to describe changing human and AI initiative during successive events in work activities. This framework is further supported by Butler et al. [6]'s observation that tasks in complex applications make diverse demands on participants. Within the framework, human-AI collaborative activities may occur in different positions, representing different degrees of human or AI control at various moments, but always with humans in the loop.

In general, AI support was desired by participants to inform missing perspectives or help people get unstuck (also described by Tholander and Jonsson [113]). It was not desired in instances where the unique perspectives and creativity of humans comprise the value of the work. Participants warned against AI use in divergent thinking, where the unique lived experiences and critical thinking skills of stakeholders and team members are indispensable. Hence, this phase of group ideation falls in the most human-initiated dimension of Figure 4. In moments when participants get stuck or certain perspectives are missing, the facilitator may choose to invoke AI support and allow the AI to contribute ideas – hence increasing AI initiative. Participants' comments indicated that there are also opportunities for AI initiative in the articulation work [97] that takes place prior to and after an ideation session, such as writing workshop questions and summaries. They also envisioned that the AI could proactively surface metadata for the facilitator to consider during the workshop, such as nudges to stay on time.

In the current configuration of the *Collaborative Canvas*, the AI is solely reactive – it only generates content upon a user's request.

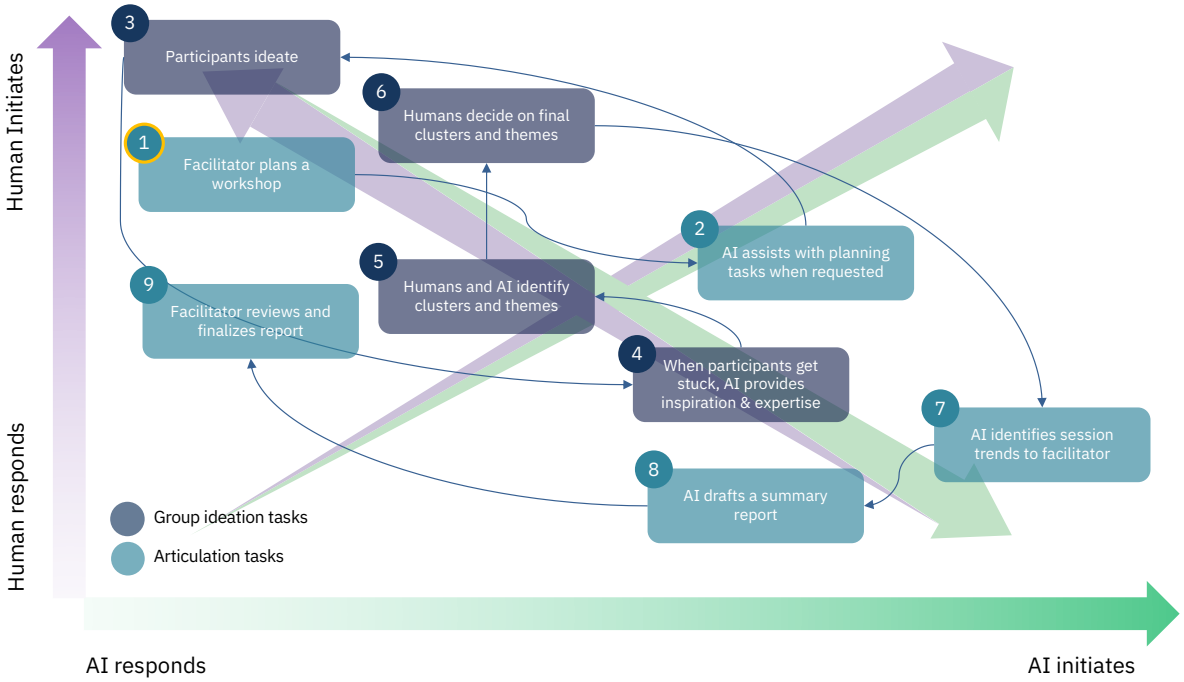


Figure 4: Desired human and AI initiative in a group ideation workflow. Tasks towards the upper left lean more towards human initiative, with divergent thinking as the task where participants wanted the least AI intervention. Tasks towards the bottom right lean more towards AI initiative, with selected types of AI summarization and analysis as opportunities for AI support. In this workflow, tasks fall along Parasuraman et al. [83]'s trade-off dimension (represented by the arrow in the background spanning from upper left to lower right), in which more human initiative means less AI initiative, and vice versa. We also make the distinction between work that occurs during a group ideation session (in dark blue) and the invisible articulation work [97] that occurs behind the scenes to enable the session (in lighter blue).

This behavior resonated well with the desire to limit AI intervention during divergent ideation. At the same time, participants' desire for more active AI support with facilitation tasks indicates opportunities to introduce proactivity in automating those tasks. In line with Maier et al. [71]'s findings on preference for human facilitators over AI and Tholander and Jonsson [113]'s report on the importance of human sense-making, proactivity should be incorporated in a way that *augments* facilitators by keeping humans in the loop and allowing them to retain control. For example, a facilitator may configure proactivity based on their personal preference or the task at hand, including the degree of proactivity at different stages of ideation, whether to allow asynchronous turn-taking, and where proactivity is surfaced (e.g. in a private vs. shared space, in sticky notes vs. a chat interface).

The AI may be able to offer more informed proactive support if given context on group members and their ideas. For example, awareness of content added to the shared canvas could help the AI generate content that is more aligned with participants' desired context and level of detail and avoid replicating existing ideas. If provided with users' backgrounds and expertise [10, 117], the AI could use this information as a facilitation aid, such as by nudging users to share ideas based on their unique perspectives. Building on our finding that the value of collaborative ideation often comes from the *process* rather than outcomes, facilitators could also call on a more proactive configuration of the AI to converse with users to explain and refine shared ideas. However, more research is needed to understand the ethical considerations, desirability, and usefulness of such features, both for facilitators and group members.

6.2 Implications for the future of AI-supported collaborative work

Our last research question focused on implications stemming from the infusion of generative AI into shared digital tools. Although our study focused on group ideation in a virtual canvas as a case study, many of these ideas are also applicable to other types of shared tools and broader domains of future collaborative work.

Selective use. Given participants' assessments of specific instances when AI can be an asset vs. when it can be a liability, users may require controls for enabling or limiting AI usage. For example, facilitators may choose to only enable generative capabilities during lulls in divergent thinking or for clustering support in convergent thinking. This approach would allow them to enhance group creativity only when needed and rely on people's innate creativity the remainder of the time. The need for facilitator-specific controls has also been present in non-AI tools; for example, Takouachet et al. [110] also identified the need for facilitator configuration in non-AI virtual canvases, including the ability to control the visibility of

sticky notes and the canvas functionalities available to participants at different times. When leveraging idea-generation capabilities, users may benefit from additional controls over the diversity of outputs. One way to address this need is by surfacing a temperature control, as proposed by Lavrič and Škraba [63], which would allow users to set a level of unexpectedness in outputs that provides novelty while remaining on-topic [112]. The model itself can also be prompted to evaluate ideas for relevance, with the less relevant ones hidden by a filter set by the user. The need for selective AI support builds on prior work that discusses users' perceptions of favorable and unfavorable AI roles [56, 102, 125].

Transparency of content ownership. Most participants found it important and more ethical to be transparent in identifying AI-generated content to the group, building on needs to disclose AI contributions in individual work [91, 111] and signify the role of the AI [118]. In collaborative, co-creative contexts, attributing AI outputs may also involve crediting the user who wrote the prompt and those who iterated on an idea – showing a version history of AI and user contributions can help groups understand where an idea came from and the role of AI within its content. However, nuances in the extent of human vs. AI ownership can impact these preferences. Lehmann et al. [64] identified a strong sense of human ownership in artifacts co-created with AI. Rezwana and Maher [94] found that ethical stances of ownership depend on the role framing of the AI. We learned that the specificity of a user's prompt and the impact of their edits to AI-generated content may impact their perceptions of how much of an idea they own. As AI-generated artifacts become more common in professional settings, these nuances can raise not only raise ethical issues, but also legal repercussions on intellectual property rights [25].

Changing needs for private digital spaces. As generative AI enters collaborative work environments, users will need ways to protect their reputation and work from AI "misbehaviors." Protections may include mechanisms to review outputs privately, modify them, and select which ones to add to a shared workspace. In studies of single-user interactions, this kind of boundary was considered useful for personal quality control [23, 42], such as programmers who selectively decided which AI-generated code to incorporate to avoid introducing problems to their code [95]. In collaborative work, designing these forms of friction [17] at key decision-making points can reduce overreliance on the AI [118] and support users with impression management [33, 42].

Towards specialized AI support. In considering ways that future generative AI capabilities could augment group work, a consistent desire was AI support that is specific to users' roles and context of work. This specificity can be achieved by providing the AI with source documents and using retrieval-augmented generation [65] to produce faithful, contextually relevant ideas, by fine-tuning the model with domain knowledge [38], or by supplying context in the prompt via few-shot learning [52]. Further research is needed to understand how context-grounded AI knowledge should be surfaced to users: as agents, as libraries, in a UI chat, or other modalities.

6.3 Positionality

We conducted this work as researchers of a large technology company; all participants were also employees of this company. Other organizations may have different collaboration and ideation practices. Due to constraints of scheduling multiple participants and researchers in a session, all participants were located in North America, with the majority in the United States, and all spoke English. This work was conducted as part of a charter to research "how to design the interactions between a SME and a generative AI system to achieve superior outcomes." We evaluated outcomes in three ways: artifact-centric (were better artifacts produced?), process-centric (how does work compare with and without AI support?), and human-centric (how do people feel while working with AI?). Our work on the *Collaborative Canvas* in this study focused on the third evaluation.

6.4 Limitations and future directions

Although we sought to understand participants' desired interactions with the AI, the study context and novelty of the tool may have influenced their use of it during the study. Additional research that embeds such generative AI capabilities within workers' tasks and projects is needed to better understand their desired AI usage in an everyday collaborative work context. Furthermore, the scope of our study was limited to participants' impressions of AI integration. Additional research that compares collaborative work with and without AI would better elucidate how its integration impacts different aspects of group dynamics, such as creativity and engagement levels.

7 CONCLUSION

The introduction of generative AI into multi-user tools raises novel considerations for the future of collaborative work. We sought to understand the unique impacts and user needs that arise from this shift by focusing on group ideation as a case study of collaborative work in a shared digital workspace. Using *Collaborative Canvas* as a probe, we conducted a study with 17 professionals experienced in facilitating or partaking in group ideation workshops. Participants recognized the potential value of AI support in group ideation, particularly in instances when certain perspectives are missing, when there is a lull in ideas, or when facilitation assistance is desired. They also identified instances when AI is detrimental – particularly in tasks where value is derived from human experiences and collaboration, and in situations where AI misbehaviors may pose a risk to their reputation. In cases where AI support is appropriate, participants' insights led us to identify unique user needs that arise from the use of AI in shared settings, beyond group ideation, including transparency of content ownership, private digital spaces, and specialized AI capabilities.

As generative AI capabilities become more advanced, we recognize an opportunity to go beyond mere automation of workflows and towards *enhancement* of work practices through tailored AI contributions. However, this vision can only be achieved by being mindful of risks such as de-skilling and loss of human collaboration, and by working towards augmentation rather than replacement of unique human processes. Our work contributes to the growing body of research aimed at thoughtfully steering the future of AI-supported collaboration with human needs at the forefront.

REFERENCES

[1] Lauren B Anderson, Dhiraj Kanneganti, Mary Bentley Houk, Rochelle H Holm, and Ted Smith. 2023. Generative AI as a tool for environmental health research translation. *Geohealth* 7, 7 (2023), e2023GH000875.

[2] Hugh Beyer and Karen Holtzblatt. 1998. Contextual Inquiry and Design.

[3] Sebastian Gregor Bouschery, Vera Blazevic, and Frank T Piller. 2023. AI-Augmented Creativity: Overcoming the Productivity Loss in Brainstorming Groups. In *Academy of Management Proceedings*, Vol. 2023. Academy of Management Briarcliff Manor, NY 10510, 11938.

[4] Virginia Braun and Victoria Clarke. 2021. Can I use TA? Should I use TA? Should I not use TA? Comparing reflexive thematic analysis and other pattern-based qualitative analytic approaches. *Counselling and psychotherapy research* 21, 1 (2021), 37–47.

[5] Shannon Briggs, Matthew Peveler, Jaimie Drozdal, Lilit Balagyozyan, Jonas Braasch, and Hui Su. 2020. Brainstorming for sensemaking in a multimodal, multiuser cognitive environment. In *HCI International 2020-Late Breaking Papers: User Experience Design and Case Studies: 22nd HCI International Conference, HCII 2020, Copenhagen, Denmark, July 19–24, 2020, Proceedings 22*. Springer, 66–83.

[6] J. Butler, S. Jaffe, N. Baym, M. Czerwinski, S. Iqbal, K. Nowak, R. Rintel, A. Sellen, M. Vorvoreanu, B. Hecht, and J. Teevan (Eds.). 2023. Microsoft New Future of Work Report 2023. *Microsoft Research Tech Report MSR- TR-2023-34* (2023). https://aka.ms/nfw2023

[7] Kevin Byron. 2012. Creative reflections on brainstorming. *London Review of Education* (2012).

[8] Marie-Laure Cahier and Suzy Canivenc. 2021. *Is Remote Working Shaping the Future of Work?* Vol. 35. Presses des Mines.

[9] Alice Cai, Steven R Rick, Jennifer L Heyman, Yanxia Zhang, Alexandre Filipowicz, Matthew Hong, Matt Klenk, and Thomas Malone. 2023. DesignAID: Using Generative AI and Semantic Diversity for Design Inspiration. In *Proceedings of The ACM Collective Intelligence Conference*. 1–11.

[10] Mustafa Mert Çelikok, Tomi Peltola, Pedram Daee, and Samuel Kaski. 2019. Interactive AI with a Theory of Mind. *arXiv preprint arXiv:1912.05284* (2019).

[11] Ting-Ju Chen and Vinayak R Krishnamurthy. 2020. Investigating a mixed-initiative workflow for digital mind-mapping. *Journal of Mechanical Design* 142, 10 (2020), 101404.

[12] Ting-Ju Chen, Ronak R Mohanty, Miguel A Hoffmann Rodriguez, and Vinayak R Krishnamurthy. 2019. Collaborative mind-mapping: a study of patterns, strategies, and evolution of maps created by peer-pairs. In *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, Vol. 59278. American Society of Mechanical Engineers, V007T06A039.

[13] Bo T Christensen and Sille Julie J Abildgaard. 2021. Kinds of 'moving'in designing with sticky notes. *Design Studies* 76 (2021), 101036.

[14] Bo T Christensen and Morten Friis-Olivarius. 2020. How do initial ideas evolve into final ones? Exploring the cognitive size, structure and life of ideas using sticky notes. In *Sticky creativity*. Elsevier, 53–75.

[15] Bo T Christensen, Kim Halskov, and Clemens N Klokmose. 2019. *Sticky Creativity: Post-it® Note Cognition, Computers, and Design*. Academic Press.

[16] Jeff Conklin and Michael L Begeman. 1988. gIBIS: A hypertext tool for exploratory policy discussion. *ACM Transactions on Information Systems (TOIS)* 6, 4 (1988), 303–331.

[17] Anna L Cox, Sandy JJ Gould, Marta E Cecchinato, Ioanna Iacovides, and Ian Renfree. 2016. Design frictions for mindful interactions: The case for microboundaries. In *Proceedings of the 2016 CHI conference extended abstracts on human factors in computing systems*. 1389–1397.

[18] Izabel Cvetkovic, Maren Gierlich-Joas, Navid Tavanapour, Nicole Debowski-Weimann, and Eva AC Bittner. 2023. Augmented Facilitation: Designing a multi-modal Conversational Agent for Group Ideation. (2023).

[19] Peter Dalsgaard, Kim Halskov, and Clemens Nylandsted Klokmose. 2020. A study of a digital sticky note design environment. In *Sticky Creativity*. Elsevier, 155–174.

[20] Rouxelle de Villiers. 2022. Design Thinking as a Problem Solving Tool. In *The Handbook of Creativity & Innovation in Business: A Comprehensive Toolkit of Theory and Practice for Developing Creative Thinking Skills*. Springer, 223–242.

[21] Melanie Duckert, Eve Hoggan, Louise Barkhuus, Pernille Bjørn, Nina Boulus-Rodje, Susanne Bødker, Naja Holten Møller, and Irina Shklovski. 2022. Work of the Future. In *Adjunct Proceedings of the 2022 Nordic Human-Computer Interaction Conference*. 1–4.

[22] Karen Leggett Dugosh, Paul B Paulus, Evelyn J Roland, and Huei-Chuan Yang. 2000. Cognitive stimulation in brainstorming. *Journal of personality and social psychology* 79, 5 (2000), 722.

[23] Camille G Endacott and Paul M Leonardi. 2022. Artificial intelligence and impression management: Consequences of autonomous conversational agents communicating on one's behalf. *Human Communication Research* 48, 3 (2022), 462–490.

[24] K Erlich, Steven Rohall, Steve Ross, Dan Gruen, Tristan Ratchford, John Patterson, Hironobu Takagi, Tatsuya Ishihara, and A Kosugo. 2012. Beyond discussions: Designing for sociability and structure. In *Proc. CSCW'12 Workshop on Collective Intelligence as Community Discourse and Action. Retrievable at: http://events. kmi. open. ac. uk/cscw-ci2012/programme-papers-demos*.

[25] Jason K Eshraghian. 2020. Human ownership of artificial creativity. *Nature Machine Intelligence* 2, 3 (2020), 157–160.

[26] Robert G Farrell, Jonathan Lenchner, Jeffrey O Kephjart, Alan M Webb, MIchael J Muller, Thomas D Erikson, David O Melville, Rachel KE Bellamy, Daniel M Gruen, Jonathan H Connell, et al. 2016. Symbiotic cognitive computing. *AI Magazine* 37, 3 (2016), 81–93.

[27] Andreas Fink, Karl Koschutnig, Mathias Benedek, Gernot Reishofer, Anja Ischebeck, Elisabeth M Weiss, and Franz Ebner. 2012. Stimulating creativity via the exposure to other people's ideas. *Human brain mapping* 33, 11 (2012), 2603–2610.

[28] Paul M Fitts, MS Viteles, NL Barr, DR Brimhall, Glen Finch, Eric Gardner, WF Grether, WE Kellum, and SS Stevens. 1951. *Human engineering for an effective air-navigation and traffic-control system, and appendixes 1 thru 3*. Technical Report. Ohio State Univ Research Foundation Columbus.

[29] Jonas Frich, Midas Nouwens, Kim Halskov, and Peter Dalsgaard. 2021. How digital tools impact convergent and divergent thinking in design ideation. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–11.

[30] Amelia Wen Jiun Gan, Quoc Dang, Blaine Western, et al. [n. d.]. AI-Mediated Group Ideation. ([n. d.]).

[31] Julia Geerts, Jan de Wit, and Alwin de Rooij. 2021. Brainstorming With a Social Robot Facilitator: Better Than Human Facilitation Due to Reduced Evaluation Apprehension? *Frontiers in Robotics and AI* 8 (2021), 657291.

[32] Katy Ilonka Gero. 2023. *AI and the Writer: How Language Models Support Creative Writers*. Columbia University.

[33] Erving Goffman et al. 1959/2002. The presentation of self in everyday life. 1959. *Garden City, NY* 259 (1959/2002).

[34] Sandy JJ Gould, Anna Rudnicka, Dave Cook, Marta E Cecchinato, Joseph W Newbold, and Anna L Cox. 2023. Remote work, work measurement and the state of work research in human-centred computing. *Interacting with Computers* (2023), iwad014.

[35] Matthias Griebel, Christoph Flath, and Sascha Friesike. 2020. Augmented creativity: Leveraging artificial intelligence for idea generation in the creative sphere. (2020). https://aisel.aisnet.org/ecis2020_rip

[36] Raja Gumienny, Lutz Gericke, Matthias Wenzel, and Christoph Meinel. 2013. Supporting creative collaboration in globally distributed companies. In *Proceedings of the 2013 conference on Computer supported cooperative work*. 995–1007.

[37] Raja Gumienny, Christoph Meinel, Lutz Gericke, Matthias Quasthoff, Peter LoBue, and Christian Willems. 2011. Tele-board: enabling efficient collaboration in digital design spaces across time and distance. *Design Thinking: Understand–Improve–Apply* (2011), 147–164.

[38] Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A Smith. 2020. Don't stop pretraining: Adapt language models to domains and tasks. *arXiv preprint arXiv:2004.10964* (2020).

[39] S Harvey and C Kou. 2017. Social Processes and Team Creativity. *Reiter-PalmonR.(Ed.), Team Creativity and Innovation* (2017), 87–128.

[40] Sarah Harvey and Chia-Yu Kou. 2013. Collective engagement in creative tasks: The role of evaluation in the creative process in groups. *Administrative science quarterly* 58, 3 (2013), 346–386.

[41] Brandon Harwood. 2023. CHAI-DT: A Framework for Prompting Conversational Generative AI Agents to Actively Participate in Co-Creation. arXiv:2305.03852 [cs.HC]

[42] 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.

[43] Melville Hensey. 1993. Essential tools of total quality management. *Journal of management in engineering* 9, 4 (1993), 329–339.

[44] Matt Heying. 2023. Introducing Mural's New AI-Powered Teamwork Features. https://www.mural.co/blog/announcing-mural-ai

[45] Otmar Hilliges, Lucia Terrenghi, Sebastian Boring, David Kim, Hendrik Richter, and Andreas Butz. 2007. Designing for collaborative creative problem solving. In *Proceedings of the 6th ACM SIGCHI conference on Creativity & cognition*. 137–146.

[46] Xavier Ho and Martin Tomitsch. 2019. Affordances of brainstorming toolkits and their use in game jams. In *Proceedings of the 14th international conference on the foundations of digital games*. 1–10.

[47] Karen Holtzblatt. 2007. Contextual design. In *The human-computer interaction handbook*. CRC press, 975–990.

[48] Eric Horvitz. 1999. Principles of Mixed-Initiative User Interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Pittsburgh, Pennsylvania, USA) *(CHI '99)*. Association for Computing Machinery, New York, NY, USA, 159–166. https://doi.org/10.1145/302979.303030

[49] Hilary Hutchinson, Wendy Mackay, Bo Westerlund, Benjamin B Bederson, Allison Druin, Catherine Plaisant, Michel Beaudouin-Lafon, Stéphane Conversy, Helen Evans, Heiko Hansen, et al. 2003. Technology probes: inspiring design for and with families. In *Proceedings of the SIGCHI conference on Human factors*

in computing systems. 17–24.
[50] Mads Møller Jensen, Sarah-Kristin Thiel, Eve Hoggan, and Susanne Bødker. 2018. Physical versus digital sticky notes in collaborative ideation. *Computer Supported Cooperative Work (CSCW)* 27 (2018), 609–645.
[51] Birgit Helene Jevnaker and Johan Leif Olaisen. 2022. The Knowledge work of the Future and the Future of Knowledge work. In *Proceedings of the 23rd European Conference on Knowledge Management, Vol. 23 No. 1 (2022)*. aci.
[52] Xisen Jin, Bill Yuchen Lin, Mohammad Rostami, and Xiang Ren. 2021. Learn continually, generalize rapidly: Lifelong knowledge accumulation for few-shot learning. *arXiv preprint arXiv:2104.08808* (2021).
[53] Vera John-Steiner. 2000. *Creative collaboration*. Oxford University Press.
[54] J Joosten, V Bilgram, A Hahn, and D Totzek. 2024. Comparing the Ideation Quality of Humans With Generative Artificial Intelligence. *IEEE Engineering Management Review* (2024).
[55] Pegah Karimi, Jeba Rezwana, Safat Siddiqui, Mary Lou Maher, and Nasrin Dehbozorgi. 2020. Creative sketching partner: an analysis of human-AI co-creativity. In *Proceedings of the 25th International Conference on Intelligent User Interfaces*. 221–230.
[56] Jihyun Kim, Kelly Merrill Jr, and Chad Collins. 2021. AI as a friend or assistant: The mediating role of perceived usefulness in social AI vs. functional AI. *Telematics and Informatics* 64 (2021), 101694.
[57] A Baki Kocaballi. 2023. Conversational ai-powered design: Chatgpt as designer, user, and product. *arXiv preprint arXiv:2302.07406* (2023).
[58] Janin Koch, Andrés Lucero, Lena Hegemann, and Antti Oulasvirta. 2019. May AI? Design ideation with cooperative contextual bandits. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–12.
[59] Christian Kohls. 2015. The Magic 5 of Innovation. In *Writers' Workshop of 22nd Conference on Pattern Lan-guages of Programs*.
[60] Anton Korinek. 2023. Generative AI for economic research: Use cases and implications for economists. *Journal of Economic Literature* 61, 4 (2023), 1281–1317.
[61] David Kut'ák, Milan Doležal, Bojan Kerous, Zdenek Eichler, Jiří Vašek, and Fotis Liarokapis. 2019. An interactive and multimodal virtual mind map for future workplace. *Frontiers in ICT* 6 (2019), 14.
[62] Lauramaria Laine. 2021. *Ideating With an Artificial Colleague: Exploring Advertising Creatives' Attitudes Towards Human-AI Collaboration*. Master's thesis.
[63] Franc Lavrič and Andrej Škraba. 2023. Brainstorming Will Never Be the Same Again—A Human Group Supported by Artificial Intelligence. *Machine Learning and Knowledge Extraction* 5, 4 (2023), 1282–1301.
[64] Florian Lehmann, Niklas Markert, Hai Dang, and Daniel Buschek. 2022. Suggestion lists vs. continuous generation: Interaction design for writing with generative models on mobile devices affect text length, wording and perceived authorship. In *Proceedings of Mensch und Computer 2022*. 192–208.
[65] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems* 33 (2020), 9459–9474.
[66] Lassi A Liikkanen, Kai Kuikkaniemi, Petri Lievonen, and Pauli Ojala. 2011. Next step in electronic brainstorming: collaborative creativity with the web. In *CHI'11 Extended Abstracts on Human Factors in Computing Systems*. 2029–2034.
[67] Yuyu Lin, Jiahao Guo, Yang Chen, Cheng Yao, and Fangtian Ying. 2020. It is your turn: Collaborative ideation with a co-creative robot through sketch. In *Proceedings of the 2020 CHI conference on human factors in computing systems*. 1–14.
[68] Andrés Lucero. 2015. Using affinity diagrams to evaluate interactive prototypes. In *Human-Computer Interaction–INTERACT 2015: 15th IFIP TC 13 International Conference, Bamberg, Germany, September 14-18, 2015, Proceedings, Part II 15*. Springer, 231–248.
[69] Wayne G Lutters. 2004. Identifying expertise markers in routine work artifacts: exploring micronote taking behavior. *Computers in human behavior* 20, 2 (2004), 185–199.
[70] Allan Maclean and Diane McKerlie. 1995. Design space analysis and use representations. *Scenario-based design: envisioning work and technology in system development* (1995), 183–207.
[71] Torsten Maier, Nicolas F Soria Zurita, Elizabeth Starkey, Daniel Spillane, Christopher McComb, and Jessica Menold. 2022. Comparing human and cognitive assistant facilitated brainstorming sessions. *Journal of Engineering Design* 33, 4 (2022), 259–283.
[72] Rajshri Mainthia, Timothy Lockney, Alexandr Zotov, Daniel J France, Marc Bennett, Paul J St Jacques, William Furman, Stephanie Randa, Nancye Feistritzer, Roland Eavey, et al. 2012. Novel use of electronic whiteboard in the operating room increases surgical team compliance with pre-incision safety practices. *Surgery* 151, 5 (2012), 660–666.
[73] Christopher McComb, Peter Boatwright, and Jonathan Cagan. 2023. FOCUS AND MODALITY: DEFINING A ROADMAP TO FUTURE AI-HUMAN TEAMING IN DESIGN. *Proceedings of the Design Society* 3 (2023), 1905–1914.
[74] Inc. Microsoft. 2023. Welcome to Copilot in Whiteboard. https://support.microsoft.com/en-gb/office/welcome-to-copilot-in-whiteboard-17e8cddb-9bae-4813-bd2b-a9f108b0b43e
[75] Paul A Mongeau and Mary Claire Morr. 1999. Reconsidering brainstorming. *Group Facilitation* 1 (1999), 14.
[76] Michael Muller. 2023. Exploring Human-AI Co-Creativity under Human Control: Framing, Reframing, Brainstorming, and Future Challenges (workshop keynote). (2023).
[77] Michael Muller, Heloisa Candello, and Justin Weisz. 2023. Interactional Co-Creativity of Human and AI in Analogy-Based Design. In *International Conference on Computational Creativity*.
[78] Michael Muller and Justin Weisz. 2022. Frameworks for Collaborating Humans and AIs: Sequence and Sociality in Organizational Applications. In *CHIWORK*.
[79] Michael Muller and Justin Weisz. 2023. *Analogies-based design using a generative AI application: A play in three acts*. http://studiolab.ide.tudelft.nl/studiolab/genai-dis2023/files/2023/07/DIS_2023_workshop___analogies-2.pdf
[80] Isabel Munoz, Steve Sawyer, and Michael Dunn. 2022. New futures of work or continued marginalization? The rise of online freelance work and digital platforms. In *2022 Symposium on Human-Computer Interaction for Work*. 1–7.
[81] Bernard A Nijstad, Wolfgang Stroebe, and Hein FM Lodewijkx. 2003. Production blocking and idea generation: Does blocking interfere with cognitive processes? *Journal of experimental social psychology* 39, 6 (2003), 531–548.
[82] Alex F Osborn. 1953. Applied imagination. (1953).
[83] Raja Parasuraman, Thomas B Sheridan, and Christopher D Wickens. 2000. A model for types and levels of human interaction with automation. *IEEE Transactions on systems, man, and cybernetics-Part A: Systems and Humans* 30, 3 (2000), 286–297.
[84] Paul B Paulus. 2015. Electronic brainstorming research and its implications for e-planning. *International Journal of E-Planning Research (IJEPR)* 4, 1 (2015), 42–53.
[85] Paul B Paulus and Vincent R Brown. 2003. Enhancing ideational creativity in groups. *Group creativity: Innovation through collaboration* (2003), 110–136.
[86] Paul B Paulus and Mary Dzindolet. 2008. Social influence, creativity and innovation. *Social Influence* 3, 4 (2008), 228–247.
[87] Paul B Paulus, Mary Dzindolet, and Nicholas W Kohn. 2012. Collaborative creativity—Group creativity and team innovation. In *Handbook of organizational creativity*. Elsevier, 327–357.
[88] Paul B Paulus, Toshihiko Nakui, and Vicky L Putman. 2006. Group brainstorming and teamwork: Some rules for the road to innovation. *Creativity and innovation in organizational teams* (2006), 69–86.
[89] Matthew Peveler, Jeramey Tyler, Dacoda B Nelson, Renato Cerqueira, and Hui Su. 2020. Browser Based Digital Sticky Notes for Design Thinking. In *Companion Publication of the 2020 ACM Designing Interactive Systems Conference*. 349–352.
[90] Alain Pinsonneault, Henri Barki, R Brent Gallupe, and Norberto Hoppen. 1999. Electronic brainstorming: The illusion of productivity. *Information Systems Research* 10, 2 (1999), 110–133.
[91] Sebastian Porsdam Mann, Brian D Earp, Sven Nyholm, John Danaher, Nikolaj Møller, Hilary Bowman-Smart, Joshua Hatherley, Julian Koplin, Monika Plozza, Daniel Rodger, et al. 2023. Generative AI entails a credit–blame asymmetry. *Nature Machine Intelligence* (2023), 1–4.
[92] Adam Procter. 2023. *Project nodenogg. in*. Ph. D. Dissertation. University of Southampton.
[93] Inc. Realtime Board. 2023. Miro announces AI-powered solution to supercharge innovation for global enterprises. https://miro.com/newsroom/miro-announces-ai-powered-solution-to-supercharge-innovation-for-global-enterprises/
[94] Jeba Rezwana and Mary Lou Maher. 2023. Designing creative AI partners with COFI: A framework for modeling interaction in human-AI co-creative systems. *ACM Transactions on Computer-Human Interaction* 30, 5 (2023), 1–28.
[95] Steven I Ross, Fernando Martinez, Stephanie Houde, Michael Muller, and Justin D Weisz. 2023. The programmer's assistant: Conversational interaction with a large language model for software development. In *Proceedings of the 28th International Conference on Intelligent User Interfaces*. 491–514.
[96] Eric L Santanen, Robert O Briggs, and G-J De Vreede. 2000. The cognitive network model of creativity: A new causal model of creativity and a new brainstorming technique. In *Proceedings of the 33rd Annual Hawaii International Conference on System Sciences*. IEEE, 10–pp.
[97] Kjeld Schmidt. 2011. *Cooperative work and coordinative practices: Contributions to the conceptual foundations of Computer-Supported Cooperative Work (CSCW)*. Springer Science & Business Media.
[98] Tim Seitz and Tim Seitz. 2020. The Materiality of Design Thinking. *Design Thinking and the New Spirit of Capitalism: Sociological Reflections on Innovation Culture* (2020), 45–81.
[99] Thomas B Sheridan and William L Verplank. 1978. *Human and computer control of undersea teleoperators*. Technical Report. Massachusetts Inst of Tech Cambridge Man-Machine Systems Lab.
[100] Patrick C Shih, David H Nguyen, Sen H Hirano, David F Redmiles, and Gillian R Hayes. 2009. GroupMind: supporting idea generation through a collaborative

mind-mapping tool. In *Proceedings of the 2009 ACM International Conference on Supporting Group Work*. 139–148.
[101] Ben Shneiderman. 2020. Human-centered artificial intelligence: Reliable, safe & trustworthy. *International Journal of Human–Computer Interaction* 36, 6 (2020), 495–504.
[102] Ben Shneiderman. 2022. Ensuring Human Control over AI-Infused Systems. https://www.nationalacademies.org/news/2022/04/ensuring-human-control-over-ai-infused-systems
[103] David Sibbet. 2010. *Visual meetings: How graphics, sticky notes and idea mapping can transform group productivity*. John Wiley & Sons.
[104] Dominik Siemon. 2022. Elaborating team roles for artificial intelligence-based teammates in human-AI collaboration. *Group Decision and Negotiation* 31, 5 (2022), 871–912.
[105] Dominik Siemon, Edona Elshan, Triparna de Vreede, Sarah Oeste-Reiß, Gert-Jan de Vreede, and Philipp Ebel. 2022. Examining the Antecedents of Creative Collaboration with an AI Teammate. *ICIS Proceedings 2022* (2022).
[106] David Socha, Troy Frever, and Chunchao Zhang. 2015. Using a large whiteboard wall to support software development teams. In *2015 48th Hawaii International Conference on System Sciences*. IEEE, 5065–5072.
[107] Robyn Speer, Joshua Chin, and Catherine Havasi. 2017. Conceptnet 5.5: An open multilingual graph of general knowledge. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 31.
[108] Wolfgang Stroebe, Bernard A Nijstad, and Eric F Rietzschel. 2010. Beyond productivity loss in brainstorming groups: The evolution of a question. In *Advances in experimental social psychology*. Vol. 43. Elsevier, 157–203.
[109] Syduzzaman Syduzzaman, Mahbubor Rahman, Mazedul Islam, Ahashan Habib, and Ahmed Sharif. 2014. Implementing total quality management approach in garments industry. *European Scientific Journal* 10, 34 (2014).
[110] Nawel Takouachet, Jérémy Legardeur, and Iban Lizarralde. 2014. 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 (Ergo'IA '14)*. Association for Computing Machinery, New York, NY, USA, 20–23. https://doi.org/10.1145/2671470.2671473
[111] Arthur Tang, Kin-Kit Li, Kin On Kwok, Liujiao Cao, Stanley Luong, and Wilson Tam. 2023. The importance of transparency: Declaring the use of generative artificial intelligence (AI) in academic writing. *Journal of Nursing Scholarship* (2023).
[112] Pittawat Taveekitworachai and Ruck Thawonmas. 2023. Enhancing Novelty in ChatGPT Responses: Incorporating Random Word Brainstorming *(IAIT '23)*. Association for Computing Machinery, New York, NY, USA, Article 2, 7 pages. https://doi.org/10.1145/3628454.3628456
[113] Jakob Tholander and Martin Jonsson. 2023. Design Ideation with AI - Sketching, Thinking and Talking with Generative Machine Learning Models. *Proceedings of the 2023 ACM Designing Interactive Systems Conference* (2023). https://api.semanticscholar.org/CorpusID:259376359
[114] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. Llama 2: Open Foundation and Fine-Tuned Chat Models. arXiv:2307.09288 [cs.CL]
[115] Jack Tsao and Collier Nogues. 2024. Beyond the author: Artificial intelligence, creative writing and intellectual emancipation. *Poetics* 102 (2024), 101865.
[116] Davor Vuchkovski, Maja Zalaznik, Maciej Mitręga, and Gregor Pfajfar. 2023. A look at the future of work: The digital transformation of teams from conventional to virtual. *Journal of Business Research* 163 (2023), 113912.
[117] Qiaosi Wang, Koustuv Saha, Eric Gregori, David Joyner, and Ashok Goel. 2021. Towards mutual theory of mind in human-ai interaction: How language reflects what students perceive about a virtual teaching assistant. In *Proceedings of the 2021 CHI conference on human factors in computing systems*. 1–14.
[118] Justin D Weisz, Jessica He, Michael Muller, Gabriela Hoefer, Rachel Miles, and Werner Geyer. 2024. Design Principles for Generative AI Applications. *arXiv preprint arXiv:2401.14484* (2024).
[119] Justin D Weisz, Michael Muller, Stephanie Houde, John Richards, Steven I Ross, Fernando Martinez, Mayank Agarwal, and Kartik Talamadupula. 2021. Perfection not required? Human-AI partnerships in code translation. In *26th International Conference on Intelligent User Interfaces*. 402–412.
[120] Britt Wieland, Jan de Wit, and Alwin de Rooij. 2022. Electronic Brainstorming With a Chatbot Partner: A Good Idea Due to Increased Productivity and Idea Diversity. *Frontiers in Artificial Intelligence* 5 (2022), 880673.
[121] Dennis Wixon, Alicia Flanders, and Minette A Beabes. 1996. Contextual inquiry: grounding your design in user's work. In *Conference companion on Human factors in computing systems*. 354–355.
[122] Tongshuang Wu, Kenneth Koedinger, et al. 2023. Is AI the better programming partner? Human-Human Pair Programming vs. Human-AI pAIr Programming. *arXiv preprint arXiv:2306.05153* (2023).
[123] Xingjiao Wu, Luwei Xiao, Yixuan Sun, Junhang Zhang, Tianlong Ma, and Liang He. 2022. A survey of human-in-the-loop for machine learning. *Future Generation Computer Systems* 135 (2022), 364–381.
[124] Zhuohao Wu, Danwen Ji, Kaiwen Yu, Xianxu Zeng, Dingming Wu, and Mohammad Shidujaman. 2021. AI creativity and the human-AI co-creation model. In *Human-Computer Interaction. Theory, Methods and Tools: Thematic Area, HCI 2021, Held as Part of the 23rd HCI International Conference, HCII 2021, Virtual Event, July 24–29, 2021, Proceedings, Part I 23*. Springer, 171–190.
[125] Meng Yin, Shiyao Jiang, and Xiongying Niu. 2024. Can AI really help? The double-edged sword effect of AI assistant on employees' innovation behavior. *Computers in Human Behavior* 150 (2024), 107987.
[126] Chiu Yu-Han and Chen Chun-Ching. 2023. Investigating the Impact of Generative Artificial Intelligence on Brainstorming: A Preliminary Study. In *2023 International Conference on Consumer Electronics-Taiwan (ICCE-Taiwan)*. IEEE, 193–194.
[127] Rui Zhang, Wen Duan, Christopher Flathmann, Nathan McNeese, Guo Freeman, and Alyssa Williams. 2023. Investigating AI Teammate Communication Strategies and Their Impact in Human-AI Teams for Effective Teamwork. *Proceedings of the ACM on Human-Computer Interaction* 7, CSCW2 (2023), 1–31.

Received 16 February 2024