GraftMind: Facilitating Group Ideation with AI-Mediated Idea Sharing

Jiale li
The Hong Kong University of Science and Technology (Guangzhou)
Guangzhou, China
jli022@connect.hkust-gz.edu.cn

Yixuan Fang
The Hong Kong University of Science and Technology (Guangzhou)
Guangzhou, China
yfang073@connect.hkust-gz.edu.cn

Xiangyang He
The Hong Kong University of Science and Technology (Guangzhou)
Guangzhou, China
xhe943@connect.hkust-gz.edu.cn

Mingming Fan*
The Hong Kong University of Science and Technology (Guangzhou)
Guangzhou, China
The Hong Kong University of Science and Technology
Hong Kong, China
mingmingfan@ust.hk

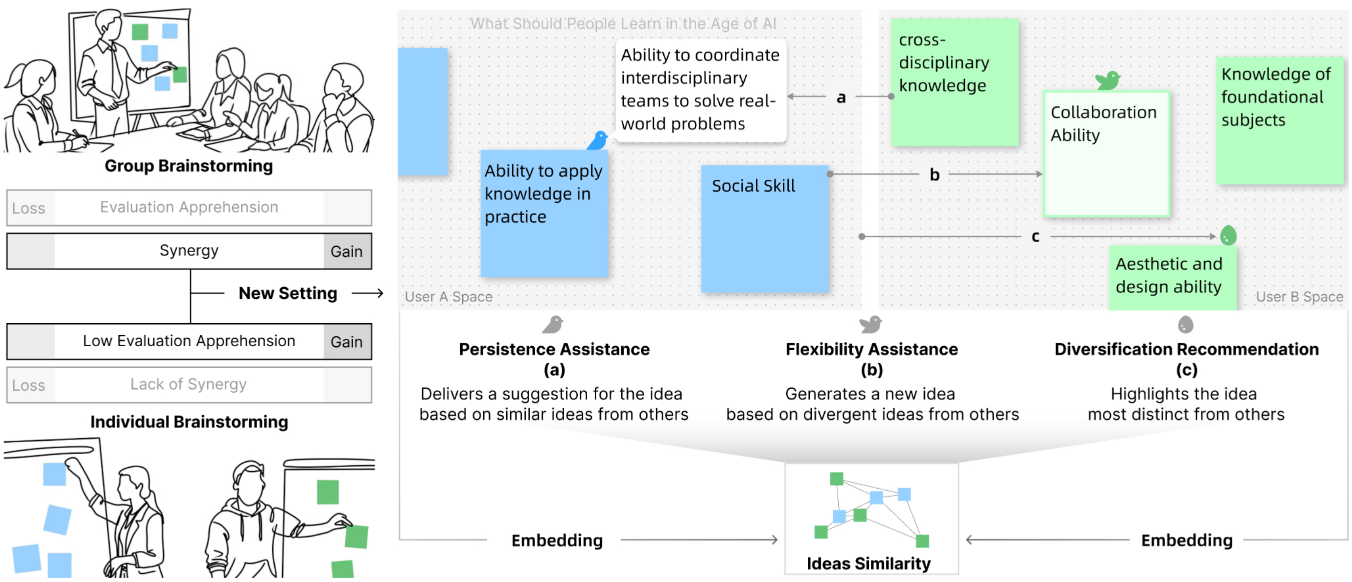

Figure 1: To integrate the advantages of group and individual brainstorming, the proactive AI mediator in *GraftMind* facilitates real-time idea sharing across private workspaces through three types of assistance. (a) Leveraging relevant ideas from other group members to provide persistence assistance. (b) Utilizing divergent ideas from others to offer flexibility assistance. (c) Recommending the most distinctive idea that differentiates the current member from the rest.

Abstract

In group ideation, whether participants should ideate collaboratively or individually remains controversial. Collaborative ideation enables synergy, whereby creativity is stimulated through inspiration from others' ideas; yet it also introduces evaluation apprehension, which can inhibit creativity due to fear of judgment. In contrast, solitary ideation mitigates evaluation apprehension but cannot foster synergy. Existing hybrid approaches have attempted to alternate between the two modes to balance their strengths, but it remains underexplored how to simultaneously integrate the advantages of both settings.

Therefore, we developed *GraftMind*, a system that enables users to ideate in private workspaces while an AI mediator proactively leverages collective ideas to provide real-time ideation assistance. The results of a user study demonstrate that *GraftMind* enhances

*Corresponding author

CHI '26, Barcelona, Spain

ACM ISBN 979-8-4007-2278-3/26/04
https://doi.org/10.1145/3772318.3791388

group ideation performance; it not only enables synergy but also alleviates evaluation apprehension. Our findings underscore the potential of this novel group ideation setting.

CCS Concepts

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

Keywords

AI-mediated collaboration, Group ideation, Human-AI co-ideation

ACM Reference Format:
Jiale li, Yixuan Fang, Xiangyang He, and Mingming Fan. 2026. GraftMind: Facilitating Group Ideation with AI-Mediated Idea Sharing. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26), April 13–17, 2026, Barcelona, Spain.* ACM, New York, NY, USA, 18 pages. https://doi.org/10.1145/3772318.3791388

1 Introduction

Group brainstorming [65] and individual brainstorming [31] are two common settings in group ideation. Group brainstorming suggests that group members ideate together in a public workspace, and each person's ideas are shared with others in real time. In contrast, individual brainstorming suggests that group members first perform solitary ideation in private workspaces and then combine their ideas.

Both group and individual brainstorming have their own strengths and weaknesses [25]. When group members ideate collaboratively, seeing others' ideas may spark new thoughts that would otherwise not have emerged; such synergy [17] can effectively enhance group creativity. However, while others' ideas can be inspiring, they can also introduce factors that impair performance, such as evaluation apprehension [9]. When ideas are shared in real time, individuals naturally anticipate how their contributions might be judged. Concerns about potential criticism may cause individuals to withhold their ideas [20], restrain thinking in novel directions [95], and conform to ideas already expressed by others [68]. To mitigate these process losses, some practices recommend that group members ideate solitarily within private workspaces. In this individual brainstorming setting, each member first generates ideas alone and then combines their ideas [78]. During solitary ideation, people can explore ideas more freely, often resulting in a larger number of ideas [86]. However, this does not necessarily improve overall idea quality, particularly in terms of originality [48]. This is because, although individual brainstorming reduces the process losses of group brainstorming, it also removes important process gains. Without interacting with others' ideas during the ideation process, individuals cannot benefit from the synergy that typically arises in group brainstorming [49].

The Human-Computer Interaction (HCI) community has been exploring hybrid approaches that could integrate the strengths of group and individual brainstorming [38]. Most existing hybrid methods adopt a similar approach in which groups alternate between group brainstorming and individual brainstorming [44]. From a broad perspective, such designs combine the advantages of both modes. However, they still confine groups to a binary choice between group and individual brainstorming at any given moment [24]. Therefore, there remains potential to explore new methods that allow the benefits of both settings to be accessed simultaneously.

In this paper, we seek to explore a novel group ideation setting that integrates the strengths of group and individual brainstorming. In recent years, a growing body of research has investigated how AI can support group ideation. Concepts such as AI ideation partners [83] and group mediators [3] have demonstrated strong potential to improve individual creativity [72] and collaborative engagement [52]. We argue that such techniques can also play a valuable role in facilitating real-time idea sharing during individual brainstorming. By doing so, the group may benefit from synergy, as in group brainstorming, while experiencing lower levels of evaluation apprehension.

Building on this motivation, we developed *GraftMind*, a system that enables users to ideate in private workspaces while an AI mediator proactively leverages collective ideas to provide real-time ideation assistance. Drawing inspiration from the dual-pathway theory of creativity stimulation [63], the system provides two types of proactive assistance: (a) Persistence assistance: The system identifies an idea from other members that is most similar to the user's current idea, then generates a deepening suggestion based on this paired relationship to support the elaboration and refinement of the idea. (b) Flexibility assistance: The system selects an idea from other members that is most distinct from the user's current ideas and synthesizes it into an alternative perspective that broadens the user's thinking. In addition, to address the practical issue of idea duplication during solitary ideation [48], *GraftMind* introduces the third type of assistance: (c) Diversification recommendation: The system detects the user's idea that differs most from those of other members, marks it with a subtle hint, and encourages the user to further explore this direction to prevent duplication. All three types of proactive assistance are conceptualized through a carrier-pigeon design metaphor and are embedded directly within the digital whiteboard interface. The interaction includes two steps: users first see a small pigeon icon, and the detailed content is revealed only after voluntarily clicking on it. This design maintains an experience consistent with conventional electronic brainstorming systems while minimizing the intrusiveness of proactive assistance.

Through *GraftMind*, we aim to examine the effectiveness of the new group ideation setting. Specifically, we investigate the following research questions:

- **RQ1: Ideation Performance.** How might *GraftMind* affect group ideation performance? Whether it could enable group members to generate ideas with higher quantity and quality?
- **RQ2: Synergy.** How might *GraftMind* affect synergy? Whether it could enable group members to gain inspiration from and build upon one another's ideas?
- **RQ3: Evaluation Apprehension.** How might *GraftMind* affect evaluation apprehension? Whether it could encourage group members to express their ideas more freely?

To address these research questions, we conducted a between-subjects user study. Participants performed group ideation using *GraftMind* or a control system that offered a shared digital whiteboard for conventional group brainstorming. A total of 60 participants (28 undergraduate students and 32 graduate students) were evenly divided into the two conditions based on academic level.

Within each condition, participants were randomly assigned to ten groups of three people. Each group then completed a 20-minute ideation task on a given topic. Note that even with the same study protocol, the control group was run later than the experimental group, so condition assignment was not fully randomized.

Our results demonstrate that *GraftMind* effectively supports group ideation. Compared to participants in the control condition, those using *GraftMind* exhibited higher levels of ideation activity, generated a greater number of ideas, and achieved higher idea quality as measured by originality. Scores on the synergy perception and evaluation apprehension scales also indicated that *GraftMind* fosters group synergy while alleviating evaluation apprehension.

In summary, this paper makes the following contributions.

- We proposed a novel group ideation setting that enables users to ideate independently while accessing AI-mediated collective ideas.
- We developed *GraftMind*, a system that enables users to ideate in private workspaces while an AI mediator proactively leverages collective ideas to provide real-time ideation assistance.
- A user study with 60 participants evaluated the effectiveness of *GraftMind*. Analysis of interaction logs and related scales indicated that this setting can enhance group ideation performance relative to conventional group brainstorming. It enables synergy while alleviating evaluation apprehension.

2 Related Work

2.1 Two Brainstorming Choices in Group Ideation

Group brainstorming is one of the most widely used approaches in group ideation [30, 74]. Its origins can be traced back to the mid-20th century, when Alex Osborn introduced the principles of freely expressing ideas and delaying judgment [65, 75]. During collaborative ideation, groups can benefit from various process gains, including synergy and social facilitation [18]. Synergy, similar to what Osborn described as piggybacking [65], is often considered the most fundamental source of group performance improvement [18]. It refers to the phenomenon in which the idea of one participant triggers new ideas in another. Through this mechanism, group members can activate relevant knowledge, gain new perspectives, and ultimately enhance creativity [62]. However, alongside these gains, group brainstorming is also susceptible to factors that impair performance. For example, evaluation apprehension [95] represents a contrasting effect of synergy. When ideating together, participants may withhold ideas because they fear negative evaluation from others [82]. This may inhibit individuals from persisting in their unique thinking [95] and ultimately undermine group creativity [20]. To alleviate evaluation apprehension, some studies have examined the impact of anonymity on idea generation [40, 70], but most have reported little improvement in group performance [18], as measured by idea quality [13] and participants' engagement [27]. Pinsonneault and Heppel [69] explain that simply hiding participants' names does not fully create a sense of anonymity, since individuals still associate themselves with the ideas they generate and remain concerned about how those ideas may be evaluated.

To avoid such process losses in group brainstorming, researchers have reconsidered having group members do individual brainstorming [33, 80]. This method allows individuals to think more freely and independently. However, this benefit comes at the cost that participants losing the opportunity to gain real-time inspiration from others' ideas [42]. Moreover, individual brainstorming cannot prevent duplication during separate work and often leads to redundant efforts on the same idea [48]. This trade-off explains why individual brainstorming often leads to a greater number of ideas [86], yet shows no significant improvement in the overall quality of group output [48].

The HCI community is exploring hybrid brainstorming methods that combine the strengths of group and individual brainstorming [38, 67]. Such approaches aim to reduce evaluation apprehension while still allowing groups to benefit from synergy. Most existing hybrid brainstorming methods follow Rossiter's guideline [78], in which groups alternate between group and individual brainstorming [44]. However, these hybrid approaches place the two modes side by side instead of integrating them into a cohesive structure. At any given moment, participants remain constrained to either group or individual brainstorming. What remains underexplored is a hybrid setting that moves beyond the traditional dichotomy. Such a setting would allow synergy to emerge as in group brainstorming, while mitigating evaluation apprehension in a manner similar to individual brainstorming.

2.2 Integrating AI in Group Ideation and Existing Challenges

An increasing body of research has investigated how AI can be integrated into group ideation processes [26, 80]. Using AI to generate ideas provides a direct way to increase group ideation performance [61, 81]. However, growing concerns about cognitive dependence on AI [45] have shifted recent research toward more supportive roles for AI. Instead of providing fully formed ideas, contemporary designs aim to stimulate users' thinking and help them generate ideas on their own [2]. The goal is not for AI to simply make group ideation faster or produce more ideas, but to activate creativity [28]. In activating group creativity, the dual-pathway theory of creativity provides a foundational framework [63, 64]. It proposes flexibility and persistence strategies: The first involves offering alternative perspectives that broaden users' thought [71]. The second involves building on users' existing directions to promote more in-depth exploration [57].

Beyond idea-oriented facilitation, some studies have explored how AI might help reduce evaluation apprehension in group brainstorming. Prior work has shown that abstracting users' ideas through AI before sharing them with the group can lower evaluation apprehension. However, this approach did not function as a catalyst that stimulated additional idea generation or offered meaningful inspiration to participants [47]. Other research suggests that adding more AI teammates may help shield individuals from interpersonal evaluation concerns [64]. However, we believe that a more meaningful question is how AI can support and strengthen human collaboration rather than substitute for it [83]. The ability to interact with other humans remains an essential skill that must be developed and trained, even in the age of AI [46].

To function as a mediator rather than a substitute, AI should account for the human–human interactions within the group [84]. Instead of merely aggregating individual ideas, AI should enable groups to enhance synergy, in which the whole becomes greater than the sum of its parts [34]. Current AI mediators mainly support either the team as a whole in synchronous collaboration [50, 61, 81] or facilitate each individual separately in asynchronous collaboration scenarios [7, 16]. These mediators typically perform functions such as task allocation, information integration, and supporting the flow of discussion [1, 89]. We believe that AI mediators can also support each individual separately in asynchronous collaboration, such as individual brainstorming. By selectively sharing ideas among team members without directly exposing the original ideas, the system can compensate for the absence of synergy in individual brainstorming while also alleviating the evaluation apprehension that often arises in group brainstorming.

To enable this form of idea sharing, the AI should first have the ability to infer users' needs. In doing so, the AI becomes aware of whether a member currently requires inspiration and what types of ideas may be helpful [23, 37, 88]. Information needs are closely associated with users' cognitive states [5], which can often be inferred from their interaction patterns and the ideas they have previously contributed [21, 73, 93, 94]. Once the need is determined, the AI could share appropriate ideas from the collective information pool [6, 16]. Conceptually, this process resembles communication in multi-agent systems, where the system determines whether and when to transmit information held by one agent to another [22]. Implementations can draw on frameworks such as Collaborative Memory [77], which organizes shared knowledge using graph structures and recommends relevant content at appropriate moments based on relational patterns.

The way AI delivers assistance is another important consideration. When an AI mediator initiates proactive interventions, it is essential to avoid presenting information that conflicts with the user's current line of thought [51]. This requirement calls for accurate timing decisions as well as interaction designs that minimize disruption [76]. For example, the system can display lightweight indicators such as a badge notification [4] that users can open when they have sufficient cognitive bandwidth. Beyond the issue of interruption, integration with existing tools is also essential. Users often adopt or abandon new intelligent features based on how smoothly they fit into established workflows [43, 66]. Electronic brainstorming tools have been developed for decades and share many common interaction patterns. The current AI features in these systems are frequently implemented through a chatbot [32, 55, 61]. Although this is a practical and widely used approach, embedding the AI mediator as a separate conversational layer can create discontinuity in the user experience. If the mediator's functions are instead presented through existing interface elements, users may experience a workflow that feels more coherent and seamless.

3 Design Considerations

Drawing from the previous discussion, we summarize three key challenges, as illustrated in Figure 2: (c1) Integrate the strengths of group and individual brainstorming; enable synergy while also mitigating evaluation apprehension. Existing research, which primarily uses alternation methodologies, fails to integrate these strengths simultaneously. (c2) Develop an AI mediator capable of sharing ideas in synchronous group ideation. Existing research primarily focuses on applying AI in group brainstorming or facilitating knowledge sharing in asynchronous collaboration. Much less attention has been given to how AI can support group members who ideate individually and simultaneously. (c3) Integrate AI features into the interaction flow of existing ideation tools while ensuring low-interruption. Existing design practices lack direct design solutions, and further design exploration is needed.

To address these three challenges, we propose four design considerations that guide the development of a system supporting this novel ideation setting.

DC1: Develop an AI mediator to share ideas while the group is engaged in individual brainstorming. Building on the above discussion, we argue that group brainstorming inevitably introduces evaluation apprehension [95]. However, the lack of synergy in individual brainstorming can be addressed through idea sharing features [85]. To this end, we aim to create private ideation workspaces that incorporate a designed mechanism to facilitate the exchange of essential ideas among individuals.

DC2: Select appropriate ideas for sharing based on the user's current ideation state. Building on our concept of an AI mediator for idea sharing, the first challenge is to identify which ideas to share with a specific user. We should develop an algorithm that infers the user's current ideation state, determining which type of idea would be most effective in activating their thinking. Specifically, we draw on the dual-pathway theory to characterize the user's idea needs, and use semantic similarity to identify relevant ideas as candidates for sharing.

DC3: Process ideas before sharing to enhance their effect in activating creativity. To enhance the effect of activating creativity, we adopt idea processing techniques such as abstraction and conceptualization [29] to transform the selected ideas for sharing. By delivering directional cues rather than specific ideas, the system encourages users to engage in their own creative thinking.

DC4: Proactively deliver assistance through a low interruption and interaction-consistent design. Since users lack visibility into others' ideas during individual ideation, it is difficult for them to know whether relevant knowledge exists elsewhere and whether they should seek system assistance. To address this, the system should provide relevant support proactively at the appropriate moments [21]. At the same time, it is important to avoid disrupting users' ongoing thought processes [32, 66]. A feasible strategy is to use mechanisms similar to unread message indicators, where the system only reveals detailed content when the user chooses to open it.

We also aim to integrate proactive assistance into the interaction flow of existing ideation tools. Electronic brainstorming systems typically use whiteboards and sticky notes as core interface elements [56]. Presenting AI-delivered ideas near the corresponding sticky notes on the digital whiteboard, rather than through a separate chat panel, may provide a more coherent and familiar user experience.

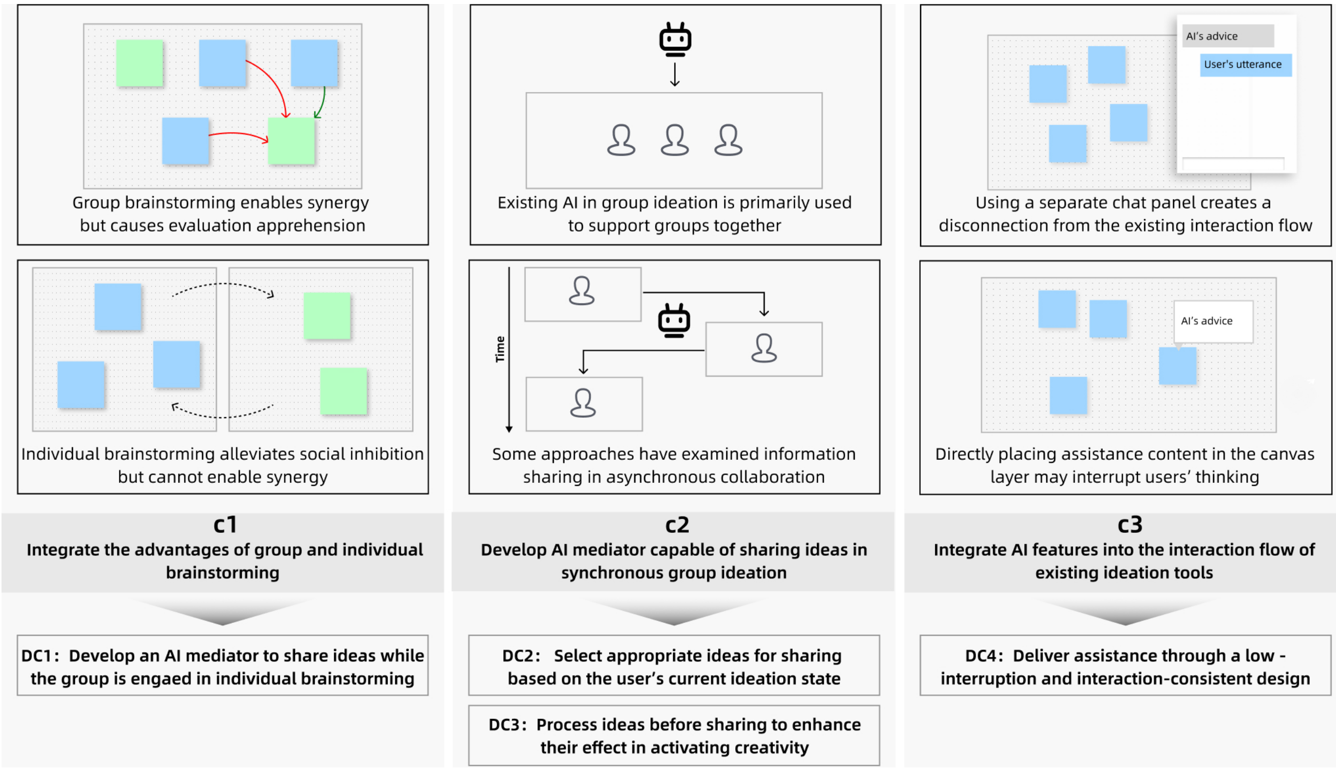

Figure 2: Three key challenges. (c1) Integrate the strengths of group and individual brainstorming. (c2) Develop an AI mediator capable of sharing ideas in synchronous group ideation. (c3) Integrate AI features into the interaction flow of existing ideation tools while ensuring low-interruption.

4 *GraftMind* System

We now discuss the implementation of *GraftMind* in response to the design considerations mentioned above. In addition to the basic functions of digital whiteboard applications for brainstorming, such as creating, editing, and dragging notes [56], the key feature of *GraftMind* is to offer ideation assistance based on collective ideas during individual brainstorming. This idea-sharing process consists of several sequential steps. First, the system infers the user's current ideation state. Second, it selects ideas from the group that match this state and generates assistance based on the selected ideas. Third, it delivers the assistance proactively with low-interruption. We now introduce each module following this temporal logic, as shown in Figure 3.

4.1 Individual Ideation State Inference

As DC2 stated, the system should first infer a user's ideation state in order to determine which idea to share next. According to the dual-pathway theory, users can benefit from two types of creativity support: flexibility and persistence. Flexibility is suitable when a user's thinking becomes overly convergent, whereas persistence is helpful when their thinking becomes excessively divergent. Therefore, the goal of state inference is to determine whether the user is currently in an overly convergent or an overly divergent state. This inference can be made by analyzing the semantic similarity among the ideas that the user has generated.

Graph structures provide a clear representation of the semantic proximity between pieces of information and allow us to identify the degree of similarity between ideas. We therefore adopt a graph-based representation to store the ideation context [39]. To construct this semantic graph, we use the Qwen3-Rerank model to compute semantic similarity between ideas. Qwen3-Rerank is among the most advanced models for estimating text similarity [90, 92]. The model outputs a similarity score ranging from 0 to 1, which we use as the edge weights in the semantic graph. Since manually setting thresholds to determine whether the user's ideation is overly divergent or convergent would be subjective, we define the user's state based on their relative internal idea similarity within the group. When a user's own idea similarity is higher than the group average, it indicates that their ideas are relatively similar, and their thinking may lack flexibility. When the similarity is lower, it indicates that their ideas may be overly divergent, making them more suitable for persistence assistance. This approach, which infers ideation states based on relative differences in thinking, also aligns with the theoretical notion of effective information complementarity [59].

The specific implementation of this state inference algorithm is as follows:

$$S_u = \frac{1}{|P(N_u)|} \sum_{(i,j)\in P(N_u)} \mathrm{sim}(i,j), \tag{1}$$

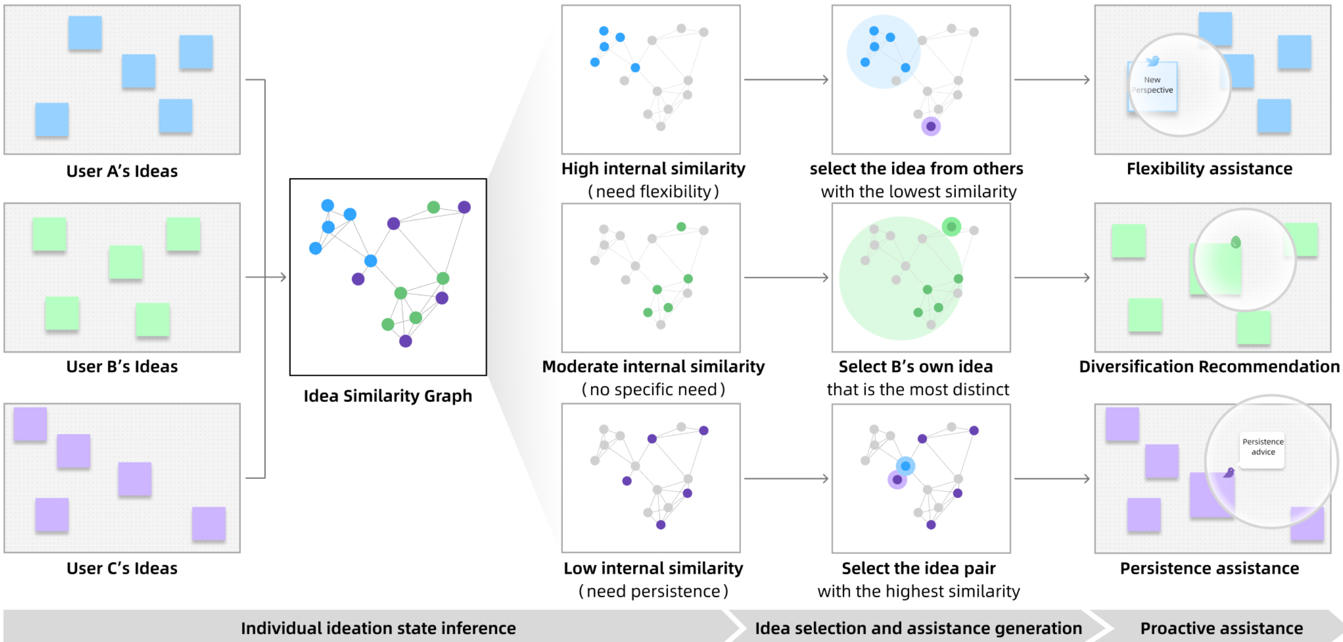

Figure 3: Idea sharing procedure of *GraftMind*. Each user's ideas are represented and stored as a semantic graph in real time. The state inference algorithm then determines the user's current cognitive needs based on the internal similarity of their ideas. The system generates assistance according to the selected strategy and finally delivers this assistance to the user proactively.

Let N_u denote the set of ideas authored by the user u, $P(N_u)$ the set of all unique note pairs, and $\mathrm{sim}(i, j)$ the semantic similarity between ideas i and j estimated by the Qwen3-Rerank model. Users with a higher internal similarity, with S_u ranked above the median, are assigned to the state that requires flexibility assistance. Those with lower internal similarity, with S_u ranked below the median, are assigned to the state that requires persistence assistance. For users whose S_u is exactly at the median, the system assumes that they can continue their current line of thinking without additional content provision. For this state, we offer a hint by recommending the most distinctive idea within their current idea set. Such diversification recommendations help prevent the duplication commonly observed in individual brainstorming.

Furthermore, to avoid repetitive user experiences, if the two most recent inferred states of a user coincide with the current one, the system substitutes the assistance strategy with an alternative. In addition, diversification recommendations are introduced stochastically across all three situations to enhance flexibility, as they are lightweight interventions applicable to all states. Through this design, the AI assistance remains both context-aware and temporally balanced.

4.2 Idea Selection and Assistance Generation

Once the intervention strategy is determined, the idea to be shared is retrieved from the semantic graph. For the user requiring persistence assistance, the system identifies the user's idea that is most strongly associated with an idea authored by others. As defined in Eq. (2), the user's idea $\hat{n}$ is the one that has the highest similarity to an external idea. This idea becomes the target for elaboration, and the associated external idea $\hat{m}$ provides the reference knowledge used to generate the persistence assistance.

$$(\hat{n}, \hat{m}) = \arg \max_{n_i \in N_u,\, m_j \in M} \mathrm{sim}(n_i, m_j). \tag{2}$$

For the user requiring flexibility assistance, the system retrieves the idea authored by other users that exhibits the greatest semantic dissimilarity relative to the ideas of the target user. Formally, for each external idea $m \in M$, we compute its minimum similarity to all ideas in N_u, as shown in Eq. (3). The idea with the lowest value of $S(m)$ is selected.

$$S(m) = \min_{n_i \in N_u} \mathrm{sim}(n_i, m). \tag{3}$$

For the user requiring diversification recommendation, the system selects the idea that is maximally distinct from other users' existing ideas and recommends it for further exploration. Formally, the most unique idea n^* is defined in Eq. (4), where $\mathrm{sim}(n_i, m_j)$ is the semantic similarity between ideas n_i and m_j, and each m_j is an idea contributed by all other members except the current user.

$$n^* = \arg \max_{n_i \in N_u} \min_{m_j \in M} \big(1 - \mathrm{sim}(n_i, m_j)\big). \tag{4}$$

Both persistence and flexibility assistance involve sharing idea content derived from other users. As highlighted in DC3, we further process these ideas to make them more effective in activating recipients' own thinking. We draw inspiration from the idea-form-issue methodology [64] and abstract each idea into a more thematic and conceptual form before sharing. The system transforms the selected idea into a higher-level thinking direction so that the shared content reflects another person's perspective without revealing their specific idea. Specifically, we use GPT-4.1, a generative model

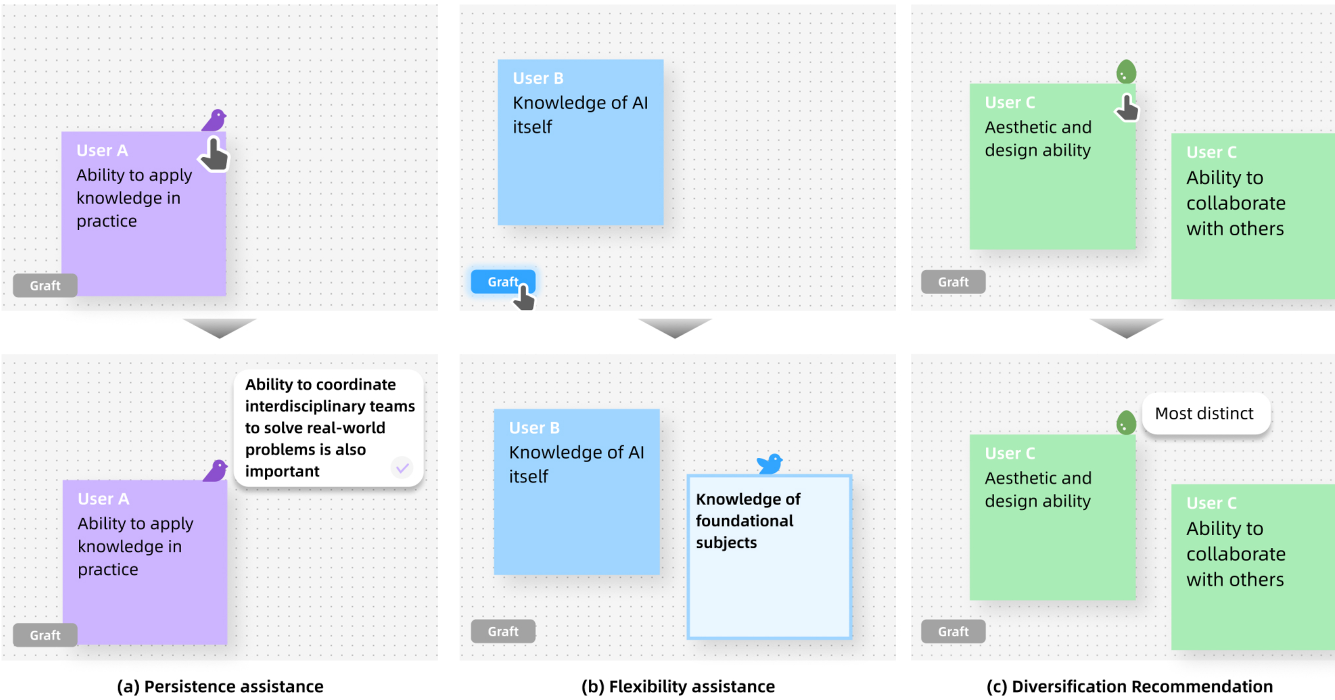

Figure 4: The interaction design of proactive assistance in *GraftMind*. (a) Persistence Assistance, (b) Flexibility Assistance, and (c) Diversification Recommendation. Each assistance type employs carrier pigeon design elements and is embedded directly into the digital whiteboard. To minimize potential disruption to users' ongoing thought processes, all proactive assistance interaction includes two steps. The system first presents a subtle notification icon and the user may choose to click the icon to reveal the detailed suggestion content.

that balances output quality and efficiency, to generate suggestions based on the selected idea. The model is instructed to summarize common themes, abstract the original content into higher-level concepts, avoid revealing or implying the original author, and generate open-ended directions that leave room for further ideation.

Some generated results are shown below. For the user who needs persistence assistance, based on the current idea, *"people should learn interdisciplinary collaboration in the AI era,"* the system infers its relevance to another user's idea, *"enhancing practical skills,"* and generates persistence assistance: *"learn to collaborate with interdisciplinary teams to solve real-world problems"*. For the user who needs flexibility assistance, when the idea *"understanding mathematics is important in the age of AI"* represents an aspect that is underrepresented in the current user's idea space, the system generates the new idea: *"knowledge of foundational disciplines"*.

4.3 Proactive Assistance

As noted in DC4, idea sharing during brainstorming should be done proactively. This requires identifying the right moment to intervene and designing low-interruption interactions that minimize disruption to the user's ongoing thought process.

GraftMind proactively provides collective-knowledge assistance without explicit user requests during the ideation phase. For the timing decision, we monitor users' interactions within the system. When a user temporarily stops ideation activities, such as creating notes, editing existing notes, or navigating the canvas, for more than 15 seconds, we infer that they may be experiencing a moment of potential blockage or a shift in attention [67]. The system then triggers ideation state inference, followed by idea selection and assistance generation. With a lightweight generation workflow, this entire process can be completed in a few seconds. Once the assistance is ready, we adopt a two-step interaction design. The system first presents a new-message indicator without revealing the detailed content, such as a notification dot, and displays the full assistance only after the user chooses to open it.

However, as mentioned in DC4, we seek to seamlessly integrate the assistance into the digital whiteboard interface. This leads us to consider how native design elements, such as notes on the canvas, could serve as carriers of assistance. The metaphor of a carrier pigeon has long been associated with the act of providing information. We found this metaphor particularly fitting for digital whiteboards, where ideas can conceptually travel across individual whiteboard spaces, carried by pigeons. In this design, a note can function both as the content that a pigeon delivers and as the place on which the pigeon temporarily perches, thereby embedding the assistance directly into the user's ongoing whiteboard-based interaction flow. This presentation can also follow a low-interruption design. When the carrier pigeon first appears on the canvas, it does not reveal its detailed suggestion, and the assistance content is shown only after the user chooses to open it.

The specific interaction designs for flexibility, persistence, and diversification recommendation assistance are illustrated in Figure 4. In persistence assistance, the pigeon perches on the note it intends to interact with, and the message expands when the user clicks on it. In flexibility assistance, the carrier pigeon flies into the current workspace carrying a new note after the user chooses to click the activated "Graft" button. For diversification recommendation, a pigeon's egg appears on the note. The pigeon's egg implies that this idea contains the potential for unique insights, suggesting to the user that, if further explored, the idea may eventually hatch into a pigeon that travels toward other workspaces.

4.4 Other Features

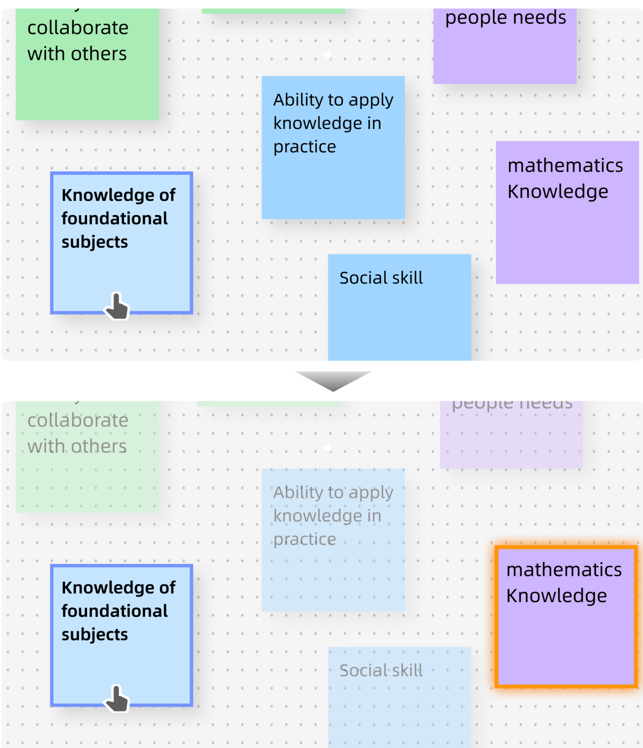

Figure 5: *GraftMind* allows users to switch to a public workspace to see which ideas their previously received assistance was based on. When users click on one of their earlier ideas, the specific idea that was referenced during its generation is highlighted in the shared space.

GraftMind also provides a chatbot to support ideation queries [15]. During the chatbot's reasoning process, we employ retrieval augmented generation [53] to enable it to draw on collective ideas as references. In addition, chatbot conversation logs are used to supplement the construction of the *GraftMind* similarity graph. Dialogs centered on a specific idea help enrich and contextualize the original idea content. It is important to note that chatbot interactions serve as an additional support feature rather than a critical or necessary component of the system. The chatbot can be collapsed and hidden if users wish to focus on their work on the digital whiteboard. *GraftMind's* idea-sharing function can still operate effectively, depending on the content of the original ideas.

GraftMind also provides a public mode in which users can view each other's ideas within a shared workspace, as shown in Figure 5. This mode allows users to see which ideas their previously received assistance was based on, which helps them assess whether the proactive assistance is explainable. This feature is primarily designed to support qualitative feedback regarding assistance explainability during the user study, while users are not permitted to switch to the public space during the ideation phase.

4.5 Implementation

During implementation, we referenced the architectures of web-based collaboration tools like Miro and Figma [60], enabling *GraftMind* to run directly in the browser and offer a user experience consistent with existing tools. The WebSocket-based Socket.IO [19] communication mechanism is used for real-time front-end synchronization. Clients subscribe to group-level events through rooms, and operations such as note creation, editing, and dragging are broadcast as event streams. When users create or modify notes, the server triggers the incremental idea graph update. In addition, we deploy a "user-pause" trigger: when the user stops interacting for a predefined duration, the system initiates state analysis and assistance generation, and the generation workflow follows a chain-structured design [87]. Given the small-scale concurrency requirements of our user study, data persistence follows a lightweight approach. Structured JSON files are written on a per-group basis, and logs are recorded in an append-only manner, removing the need for a database and facilitating rapid iteration during system development.

5 User Study

We seek to evaluate whether *GraftMind* can effectively support group ideation, enable synergy similar to group brainstorming, and simultaneously alleviate evaluation apprehension as in individual brainstorming. Therefore, we compared the performance of teams using *GraftMind* with those using a conventional group brainstorming method to examine these research questions.

5.1 Participants

We recruited 60 participants through a university mailing list. The sample comprised 28 undergraduates and 32 graduate students. All participants reported previous experience in group ideation.

5.2 Setup and Experimental Design

To avoid topic-related variability and potential learning effects within the same topic, we conducted a between-subjects study. The experimental condition used *GraftMind* to perform group ideation, while the control condition followed a conventional group brainstorming process. To support the control condition, we developed a system that replicates only the essential digital whiteboard functionality of *GraftMind*. The idea-sharing features of *GraftMind*, including proactive assistance and the chatbot, were not adopted in the control system. This control system aligns with existing group brainstorming tools such as Figma or Miro, allowing participants to collaborate in a public workspace.

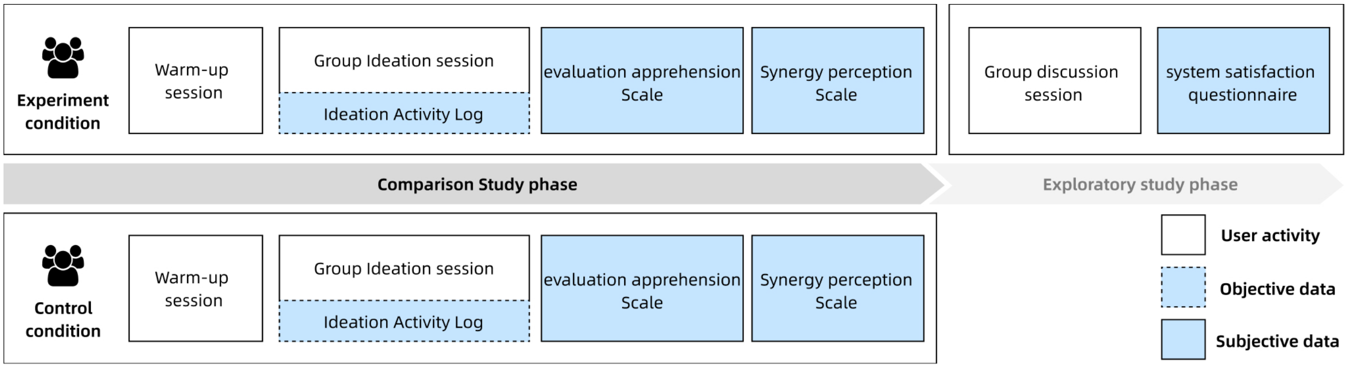

Figure 6: The procedures for both the experimental and control conditions followed the same comparison phase, which allowed us to examine differences in participants' ideation experiences and the ideas they ultimately generated. In addition, the experimental group participated in an extra exploratory phase, which enabled us to collect more qualitative feedback on the *GraftMind* system.

We assigned participants to the experimental and control conditions while ensuring a comparable distribution of education levels. Each condition included 14 undergraduate students and 16 graduate students. Within both conditions, we randomly formed 10 groups, each consisting of three participants. We selected a group size of three because it is commonly adopted in classroom discussions and project-based learning contexts. For the preparation of the group ideation topic, our aim was to ensure that the participants had similar knowledge and cognitive levels on the topic. To achieve this, we chose AI and education as the thematic boundaries, which are highly relevant to students' daily experiences. Drawing on the OECD's Future of Education report [14], we identified a well-discussed topic: What skills should students focus on cultivating for the future with the development of AI? All groups conducted brainstorming sessions on this topic.

The experimental procedure is shown in Figure 6. At the beginning, participants were introduced to the concept of brainstorming and its classic guidelines (e.g., delaying judgment and generating diverse ideas). Participants in the experimental condition were additionally introduced to the core idea-sharing mechanism of *GraftMind*, ensuring that they clearly understood how their ideas would be shared later in the study. Next, participants were given 10 minutes to freely explore the system they would later use for group ideation, ensuring they understood how to use each feature of the system. The subsequent ideation topic was not revealed during this warm-up session. Following the warm-up, participants engaged in a 20-minute brainstorming task on the designated topic. After the session, they completed scales assessing evaluation apprehension and perceived synergy. At this point, the study comparison phase was completed. The objective ideation data and subjective ratings collected in this stage allowed us to compare participants' ideation performance and user experience.

Since *GraftMind* introduces a novel ideation setting, we sought to gather more in-depth feedback on the system. To do so, we conducted an additional exploratory phase for participants in the experimental condition. In this phase, participants switched to the public workspace, where they could revisit all ideas generated by their group. We then assigned a consensus-building task that encouraged participants to actively examine related and contrasting ideas [80]. This task enabled participants to more concretely reflect on how the earlier AI-mediated assistance had influenced their ideation, for example, whether it helped them generate ideas that differed from others.

After completing this phase, participants in the experimental condition were asked to complete a system satisfaction questionnaire, rating each item on a 7-point Likert scale from 1 (very dissatisfied) to 7 (very satisfied). For items with particularly low scores (< 3) or high scores (> 5), follow-up questions were asked to identify the underlying reasons for these ratings. These interviews provided additional insights into participants' experiences with *GraftMind*.

Figure 7: A group of three participants used *GraftMind* for brainstorming. To ensure strict separation of private workspaces, the physical setup was arranged so that text displayed on other participants' screens appeared below the minimum readable size at the viewing distance.

5.3 Data Collection

The first focus of our evaluation is to examine whether *GraftMind* can support group ideation more effectively than conventional group brainstorming. Such effectiveness is reflected in whether groups generate ideas with higher quantity and quality, and whether they engage in more active ideation behaviors. While idea quantity can be measured directly by counting the number of ideas, idea quality requires an explicit definition. Following common practice in brainstorming studies, we define quality in terms of idea originality [35, 63]. We apply semantic similarity models to estimate originality and identify the number of distinctive idea clusters produced in each condition.

Ideation activity is quantified through users' interaction logs related to creative expression in the system. According to the system's interaction rules, creating a note represents generating a new idea, whereas editing a note reflects refining an existing idea. These actions correspond to two types of creative behavior: introducing new thoughts and improving existing ones. Therefore, we focus primarily on these two interaction behaviors, namely *Create Note* and *Edit Note*. *Create Note* can be directly inferred from the system's note creation operation. For *Edit Note*, users must double-click a note to enter editing mode and untick it to exit. We first calculated the number of times edit mode was activated and then performed post-processing to filter out character correction edits that did not involve semantic modification.

After establishing the basic effectiveness of the system, we further examine how the design in our experimental condition led to these differences. We recorded the first interaction behavior that participants performed after receiving each assistance, which allowed us to analyze how different assistance strategies supported ideation. In addition, during the exploratory stage, participants' scores on the satisfaction questionnaire, along with their explanations, were collected to analyze their subjective feedback on *GraftMind* features.

To analyze synergy, we designed a synergy perception scale based on the Knowledge Sharing Scale [12] to measure participants' perceptions of inspiration, knowledge exchange, and the extent to which such sharing supports their ideation process. We also inspected the ideation behaviors of participants after receiving assistance generated from other users' ideas, focusing on whether this shared information enabled them to develop new perspectives or produce more elaborate points.

For evaluation apprehension, objective behavioral indicators are limited, as it primarily reflects a psychological state. We designed an evaluation apprehension perception scale based on the Brief Fear of Negative Evaluation Scale [10], which is commonly used to assess individuals' concerns about being judged by others in collaborative settings. The items were adapted to better fit the *GraftMind* context by converting trait-based items into state-based items that capture participants' experiences during system use.

5.4 Data Analysis

Idea quantity and quality. For each condition, we computed both the total number of ideas generated per group and the number of semantically distinct idea clusters. To assess idea quality from the perspective of originality, we used embedding-based semantic clustering. All ideas were first encoded using the BGE-M3 model to obtain sentence-level embeddings [11]. We then computed pairwise cosine distances between ideas within each group and performed unsupervised clustering on the resulting distance matrix [2]. Following prior work, we adopted a density-based thresholding approach. Ideas whose embedding distances exceeded a semantic dissimilarity threshold were assigned to different clusters, resulting in a cluster structure that reflects the number of distinctive idea directions [58]. The final cluster count was regarded as the measure of group ideation quality. After obtaining idea counts and cluster counts for each group, we examined normality using the Shapiro–Wilk test and conducted one-tailed independent t-tests to compare conditions in terms of quantity and originality-based quality.

Ideation activity level. We conducted a temporal analysis that compared the average ideation activity levels of participants in the two conditions over normalized time, allowing us to examine overall differences in activity patterns between the conditions. In addition, we compared the temporal patterns of AI assistance intensity with differences in activity levels to further explore how variations in assistance intensity may have contributed to the fluctuations in user activity.

Synergy and evaluation apprehension perception scales. The responses of the synergy and evaluation apprehension perception scales were first examined for normality using the Shapiro–Wilk test and then analyzed using independent one-tailed t-tests or one-tailed Mann–Whitney U tests to assess the significance of differences between the two conditions on the two scales. Our goal was to determine whether the experimental condition achieved a comparable level of perceived synergy and a lower level of evaluation apprehension relative to the control condition.

Interviews. We performed a reflexive thematic analysis of the interview data [8] to examine how participants perceived AI assistance. Two researchers independently coded all interview transcripts. After comparing their codebooks, the researchers consolidated 76 codes with clarified definitions. These codes were then organized into 8 conceptually coherent categories and, through iterative review, grouped into 3 overarching themes that aligned with our research questions.

6 Results

We present our results regarding the three research questions of our study, which are group ideation performance, synergy facilitation, and evaluation apprehension alleviation. For each aspect, we first report the differences between the two conditions on the key measures. We then draw on the corresponding data analyzes to explain how our proposed approach contributed to these observed differences.

6.1 Group Ideation Performance

6.1.1 Performance differences based on idea quantity and quality. Shapiro–Wilk tests confirmed that the distributions of idea counts in both conditions met the normality assumption (experimental condition: $W = 0.942, p = 0.577$; control condition: $W = 0.982, p = 0.973$). We then conducted independent-samples one-tailed t-tests. As shown in Figure 8, the experimental condition produced significantly more ideas per group than the control condition ($t(58) = 2.342, p = 0.015$).

Similarly, Shapiro–Wilk tests confirmed that the distributions of idea cluster counts in both conditions met the normality assumption

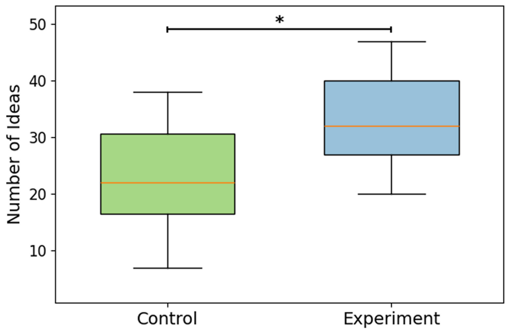

Figure 8: Comparison of idea quantity between the experimental and control groups. The experimental group generated a significantly higher number of ideas than the control group.

(experimental condition: $W = 0.922, p = 0.374$; control condition: $W = 0.934, p = 0.492$). We then conducted independent-samples one-tailed t-tests on the idea cluster counts. As shown in Figure 9, groups in the experimental condition generated significantly more idea clusters on average than those in the control condition ($t(58) = 2.177, p = 0.022$).

These results indicate that within the same amount of time, groups using *GraftMind* generated not only a larger number of ideas but also more distinctive ideas, as P3 noted, *"AI stimulates ideas that were difficult for me to think of on my own."*

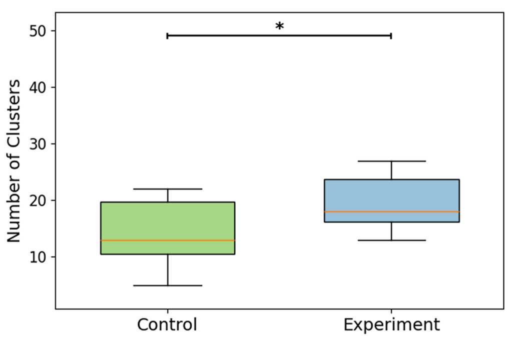

Figure 9: Comparison of idea cluster counts between the experimental and control groups. The experimental group generated a significantly higher number of idea clusters than the control group.

6.1.2 Performance differences based on activity level and its temporal pattern. We further evaluate the performance differences between the two conditions from a temporal perspective. The relationship between group ideation activity and time is shown in Figure 10. As illustrated at the top, participants in the experimental condition consistently exhibited higher ideation activity throughout the session.

The difference in activity levels between the two conditions, along with the number of AI interventions over time, is presented at the bottom. Under AI assistance, the experimental condition exhibited two additional peaks in ideation activity at normalized times 0.2 and 0.65, which correspond to approximately the 4th and 13th minutes of the 20-minute ideation session. This pattern suggests that *GraftMind* was able to reignite participants' thinking, as P13 stated, *"The generated notes give me 'aha' moments. This effectively stimulates my thinking, especially when I'm stuck."*

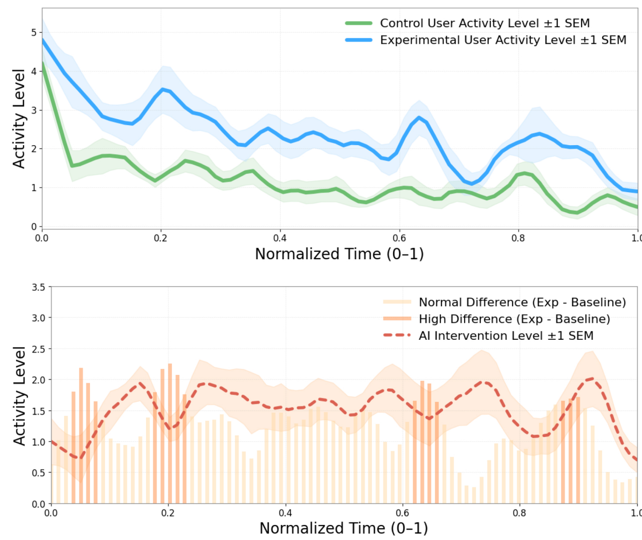

Figure 10: (Top) The change in participants' ideation activity over normalized time in both the experimental condition and the control condition. (Bottom) The difference in activity between the experimental and control conditions, along with the intensity of AI interventions, plotted over normalized time.

6.1.3 How various assistance strategies support ideation performance. Users' satisfaction with each of *GraftMind*'s main features is shown in Figure 11. Overall, most participants ($N = 27$) agreed that *GraftMind* effectively supported their brainstorming process.

The validity of the design driven by the dual-pathway theory is supported by the differences in participants' subsequent tendencies to create or edit notes following the assistance. As illustrated in Figure 12, after receiving flexibility assistance that offered a new perspective, participants were more inclined to create new notes to express new ideas. In contrast, edits to existing ideas occurred more frequently after receiving persistence assistance than after receiving flexibility assistance. This result shows that *GraftMind* supports ideation performance through both the introduction of new perspectives and support for deeper elaboration, as suggested by the dual-pathway theory.

Beyond the two strategies grounded in dual-pathway theory, the diversification recommendation received the most positive feedback ($N = 23$), as shown in Figure 11. The participants found it concise, yet informative, as P22 stated, *"The recommendation does not require me to read newly generated content, but it still provides informative cues about what other people are not thinking about."* In contrast,

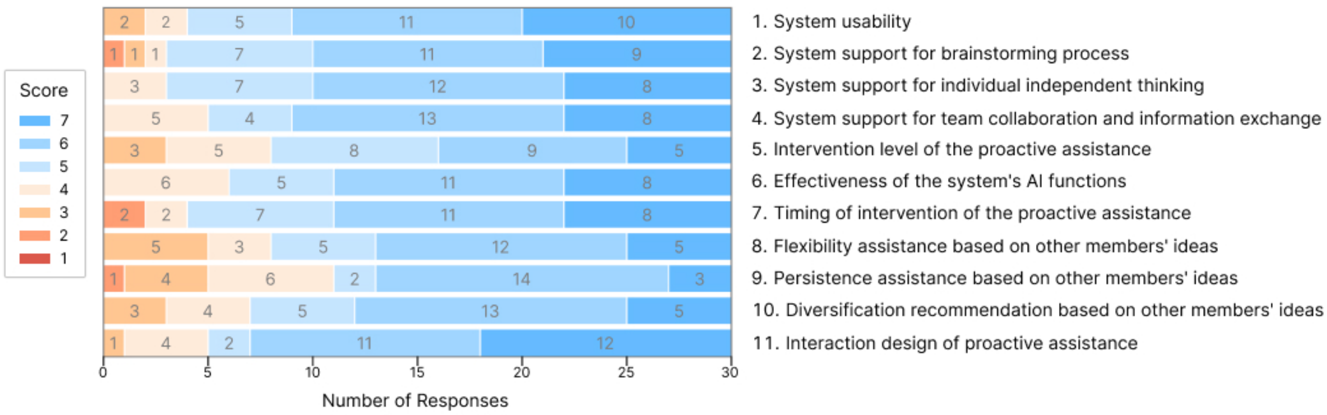

Figure 11: Participants in the experimental condition rated items on the system satisfaction questionnaire using a 7-point scale (1 = very dissatisfied, 7 = very satisfied).

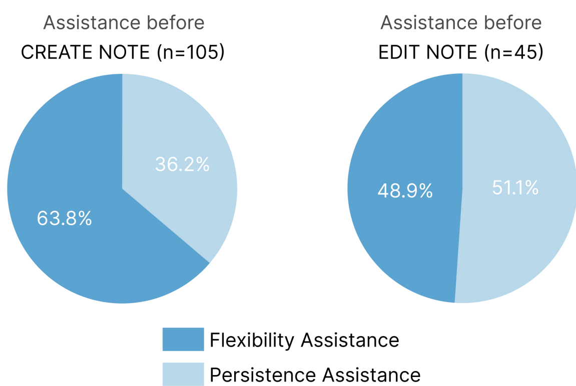

Figure 12: (Left) Distribution of assistance types preceding *Create Note.* (Right) Distribution of assistance types preceding *Edit Note. Create Note* was more prevalent following flexibility assistance, whereas *Edit Note* was more likely to occur following persistence assistance. The value of n indicates the total number of each behavior type. Because brainstorming emphasizes generating diverse ideas, *Create Note* was generally preferred over *Edit Note.*

persistence assistance received relatively lower levels of positive feedback ($N = 19$). The participants noted that they had to actively identify the relationship between the persistence assistance and the idea they targeted. As P26 explained, *"Thinking about how to apply the advice to my original idea involves some cognitive load. If I end up not using the advice, it leads to a sense of disappointment."* Preferences for flexibility assistance fell between these two types ($N = 22$). Participants reported that the workload required by flexibility assistance was moderate. They only needed to read the generated idea and decide whether to keep it, without needing to analyze relevance or re-edit existing ideas. As stated by P6, *"It directly gives me a new perspective that I would not think of myself."*

The collective idea-driven chatbot is also an idea-sharing feature in *GraftMind*. However, usage of this feature was limited: only 17 of the 30 participants in the experimental condition interacted with the chatbot, and only 4 of them engaged in more than 3 turns during the 20-minute study session. Given these limited interactions, we do not further discuss the chatbot's related feedback in this paper.

6.1.4 How the interface design supports ideation performance. Most of the participants ($N = 25$) agreed that thoughtful interaction design helped support their ideation performance. The Carrier Pigeon interface was perceived as natural, intuitive, and engaging. As P28 noted, *"I love the design of the Carrier Pigeon. It not only provides me with useful ideas but also makes the ideation process more interesting. The way it appears on the notes feels very natural and well-suited for digital whiteboards."* Moreover, the visual metaphor used for idea sharing enabled participants to intuitively perceive the flow of collective knowledge. As P10 remarked, *"The Carrier Pigeon intuitively conveyed the concept of idea sharing."*

6.2 Synergy Facilitation

6.2.1 Differences in synergy perception between the two conditions. As one of the most foundational factors supporting group ideation, the enhanced ideation performance observed in the experimental condition already provides evidence of synergy facilitation. Moreover, participants' ratings on the synergy perception scale were comparable to those in the group brainstorming condition, as shown in Figure 13. After confirming that participants' ratings in both conditions followed a normal distribution using the Shapiro–Wilk test (experimental condition: $W = 0.975, p = 0.691$; control condition: $W = 0.938, p = 0.081$). We conducted a one-tailed t-test to evaluate whether the experimental condition was significantly lower than the control condition. The result, $t(58) = -1.02, p = 0.156$, indicates that there was no significant difference between the two conditions on the synergy perception scale ratings. This suggests that participants in the experimental condition could still perceive inspiration from others' ideas even when ideating individually. As P24 remarked, *"It brings me interesting perspectives from the humanities, disciplines I am not familiar with."* Notably, there was indeed a participant with a humanities-related background in P24's group. In some cases, participants could even infer the origin of the assistance when they were familiar with each other, as P21 said, *"I could tell that this pigeon message likely came from P19 since it matched his research direction."*

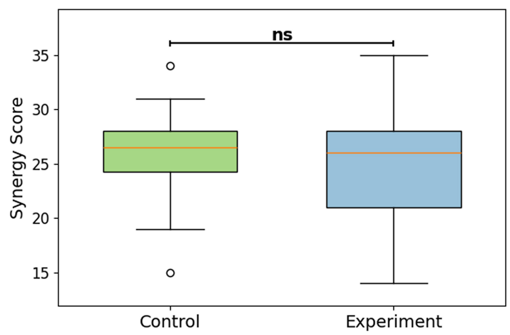

Figure 13: Comparison of synergy perception ratings between the experimental and control conditions. No significant difference was observed between the two conditions.

6.2.2 Representative cases showing synergy in the experimental condition. To further substantiate the synergy in the experimental setting, we examined the interaction logs and identified concrete instances in which participants benefited from others' ideas. Several representative cases are shown in Figure 14. These examples illustrate how participants derived inspiration through flexibility and persistence assistance and how they subsequently expressed this inspiration in their ideation process. For instance, A1 shows how the participant further elaborated on the original idea following persistence assistance, shifting from *"using AI for quick problem-solving"* to a more systematic view of *"mapping knowledge connections for deeper understanding."* In contrast, A3 demonstrated the effect of flexibility assistance. After receiving a new perspective brought by the collective knowledge retrieved from the group, the participant generated a new idea centered on *"abilities that AI cannot replace, such as proposing new design styles."*

6.2.3 How to further enhance synergy. Participants identified several factors that indicate opportunities for further refinement. Some participants suggested that providing clearer assistance provenance could enhance their sense of connection. As P21 expressed, *"If I knew which person's space the pigeon flew from, the sense of connection would feel better."* The communicative style of the AI may also influence the experience, with P20 commenting that *"The pigeon's tone felt overly rigid, like interacting with a system rather than engaging in teamwork."* These issues suggest that improving source attribution awareness and communicative warmth may further enhance synergy.

6.3 Evaluation Apprehension Alleviation

6.3.1 Differences in evaluation apprehension perception between the two conditions. Figure 15 presents participants' evaluation apprehension perception scores in the experimental and control conditions. Because the rating distributions in the two conditions did

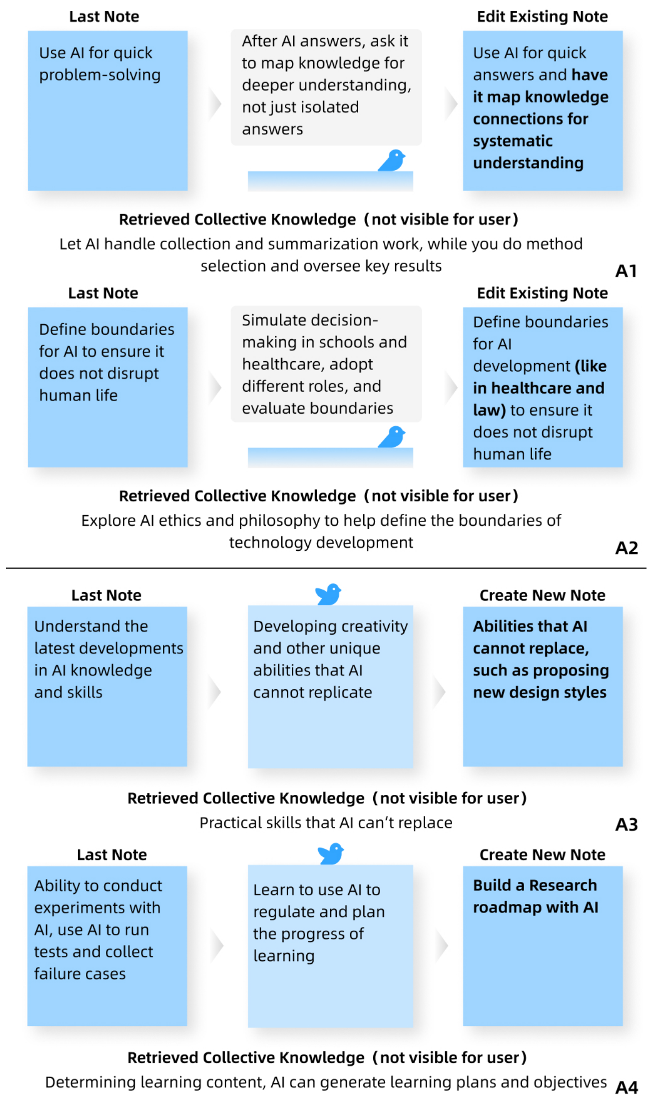

Figure 14: The changes in participants' ideation behavior before and after receiving persistence and flexibility assistance, (Top) participants enrich their existing ideas by incorporating a related idea from the persistence assistance. (Bottom) After receiving flexibility assistance, participants create ideas that are inspired by the new perspective.

not fully meet the normality assumption (experimental condition: $W = 0.885, p = 0.004$; control condition: $W = 0.944, p = 0.118$), we used a one-tailed Mann–Whitney U test to evaluate whether the experimental condition was significantly lower than the control condition. The result, $U = 141.5, p < 0.001$, indicated that the experimental condition reported significantly lower evaluation apprehension than the control condition. This suggests that *GraftMind* effectively alleviated participants' concerns about being evaluated during ideation.

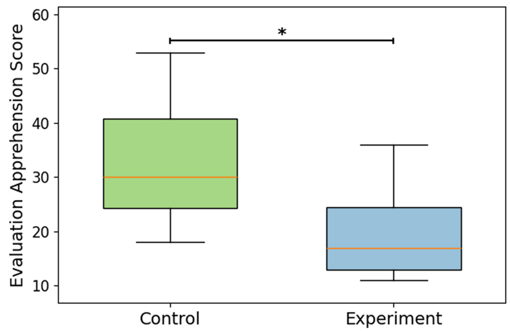

Figure 15: Comparison of evaluation apprehension perception ratings between the experimental and control conditions. Participants in the experimental condition rated significantly lower than those in the control condition.

6.3.2 How the experimental condition alleviates evaluation apprehension. The interview data further explain the findings of evaluation apprehension alleviation. Participants agree that group brainstorming often triggers self-doubt and pressure, while *GraftMind* reduces such discomfort. As P5 explained, *"In a public workspace, when I noticed that others' ideas were different from mine, I tended to interrupt my own thinking and reflect on whether I misunderstood the task. However, GraftMind encouraged me to explore more."* Likewise, P12 highlighted the reduction of conformity pressure, *"I always dwell on the ideas just generated by others in a public workspace. In contrast, I could follow my own thinking when using GraftMind."* Other participants further emphasized how the system encouraged independent exploration. P23 remarked, *"The diversification recommendation felt like an encouragement, giving me a positive cue to explore ideas different from others."* Similarly, P6 noted that *"accessing others' information in a controlled manner gives me greater autonomy to explore my own ideas."*

6.3.3 How to further enhance evaluation apprehension management. It is noteworthy that participants differed in their sensitivity to evaluation apprehension. Most participants ($N = 23$) preferred to retain the current level of sharing frequency and detail, as they were concerned that more extensive sharing might increase their anxiety. Some expressed that they were comfortable with more direct forms of idea sharing. For instance, P19 stated, *"I'm fine with the AI sharing more details with me, and I don't think I would be disturbed by such content."* Likewise, P8 expressed a desire for more direct cross-space interaction, *"When I encounter shared information that is insightful, I want to engage further with the collaborator who provided it."* These perspectives highlight that some participants with higher autonomy may benefit from more open and higher-intensity information exchange.

7 Discussion

Synthesizing the findings of our user study, we now revisit the four design considerations that guided the development of *GraftMind*. For each consideration, we discuss how the underlying design contributed to the observed effectiveness of the system and clarify the promising design spaces that deserve further exploration.

7.1 A Novel Group Ideation Setting that Integrates the Strengths of Group and Individual Brainstorming

In prior work, group ideation has largely been framed as a binary choice between group and individual brainstorming [65, 78]. Group brainstorming enables synergy but also introduces evaluation apprehension, whereas individual brainstorming alleviates evaluation apprehension but reduces the benefits of synergy. Existing hybrid approaches alternate between the two modes [44], yet they do not allow users to access the advantages of both at the same time.

GraftMind addresses this gap by introducing a new group ideation setting that allows users to work in private workspaces while an AI mediator draws on the group's collective ideas to offer real-time assistance. Based on users' generated ideas, ideation behaviors, questionnaire responses, and interview feedback, we argue that *GraftMind* effectively integrates the advantages of both group and individual brainstorming. It enables users to freely explore ideas with reduced evaluation apprehension [82, 95] while still gaining synergy from the group's collective ideas.

7.2 The Importance of Timing in Idea Sharing

The three key elements of proactive interaction are timing, content, and form [54]. Based on our analysis, we argue that timing should receive primary attention in future research. As illustrated in Figure 10, the effectiveness of AI interventions in enhancing users' ideation activity varies depending on when the intervention occurs. Participants' subjective evaluations of the assistance content were also strongly influenced by the appropriateness of its timing. When interventions were delivered during moments of cognitive blockage [67], participants tended to be more patient and more willing to explore the inspirational value of the suggestion. In contrast, when participants were already engaged in active thinking, they formed higher expectations of the content because accepting an intervention introduced additional cognitive resumption costs [41]. Taken together, group ideation requires sustained cognitive engagement, which makes participants sensitive to interruptions. As a result, their subjective impressions of an intervention are shaped by its timing even before they process the actual content, and these impressions directly influence their expectations and evaluations of the assistance itself.

7.3 Less is Better? How to Generate Assistance Based on Shared Ideas

To transform shared ideas into inspiring and open-ended forms of assistance, we designed the idea processing workflow around the dual-pathway theory. Depending on users' cognitive needs, the system generates persistence and flexibility assistance. These two forms help convert original ideas into new perspectives that stimulate divergent thinking or suggestions that encourage users to elaborate on their ongoing directions. In addition, to mitigate duplicate ideas during individual ideation, we incorporate a light hint that highlights the idea that is the most distinct, encouraging exploration of more diverse areas. Based on these results, we believe that the proposed design rationales have achieved their intended effects.

Interestingly, participants' subjective feedback indicated the highest preference for the diversification recommendation. They reported that assistance containing too much content introduced an unnecessary cognitive load. From a cognitive cost perspective, this reaction is unsurprising: when participants are required to invest more effort to interpret a suggestion, they naturally form higher expectations regarding its quality and usefulness [79]. This suggests that simplifying the content and presentation of persistence and flexibility assistance may help reduce the associated cognitive load [2]. For future work, we recommend prioritizing strategies that distill collective knowledge into concise suggestions or even non-textual hints, as these lightweight forms may better support users during cognitively demanding ideation activities.

7.4 Not Just Delivering Assistance, but Also Conveying the Concept of Idea Sharing

GraftMind's AI mediator design, embodied in the form of a "carrier pigeon," received considerable positive feedback from participants. The way pigeons paused on existing notes to deliver suggestions, left eggs to highlight unique contributions, or brought new notes into the workspace was considered consistent with the digital whiteboard interaction. This design made the overall experience feel smooth and engaging. Even without closely examining the specific content of the suggestions, participants could still perceive the concept of idea sharing through the pigeon metaphor.

The pigeon metaphor presents a promising direction for embedding AI-mediated idea sharing in a playful yet functional manner. Future work may explore how this metaphor can be extended to convey richer signals about information provenance, as suggested by several participants. Such extensions may further strengthen users' sense of working together [91]. For instance, we experimented with using pigeon colors to indicate sources; however, this approach introduced excessive color cues, cluttering the visual space. Allowing users to customize the appearance of their pigeons, while maintaining a unified color, may offer a more enjoyable and transparent way for users to understand where assistance originates.

8 Limitations and Future Work

While *GraftMind's* effectiveness in facilitating brainstorming has been initially verified, we see several limitations that need to be addressed in future research.

User study. After completing the study on *GraftMind*, we recognized the need for a control condition to further validate the system's effectiveness. Accordingly, we conducted a follow-up experiment following the same protocol to serve as the control condition. As a result, the two conditions were not fully randomized. Additionally, while our sample size of 60 participants provides preliminary insights, further research with a larger participant pool is needed to generalize and strengthen these results. Certain aspects of the experimental design could also be further improved. For example, participants in the same group could be placed in separate rooms to ensure complete isolation, thereby minimizing unintended social influence during the task.

Although participants in our current study primarily engaged with proactive assistance features, the observed effects cannot be fully attributed to this component alone. Future work should conduct ablation studies that independently evaluate other features, such as the collective idea-driven chatbot, to disentangle the respective effects of each component on ideation performance.

Detecting ideation stalls. Although participants generally responded positively to the current timing of AI assistance, several limitations remain in *GraftMind's* proactive timing mechanism. The pause-based triggers inferred from observable interaction behaviors cannot reliably capture users' underlying cognitive states when they temporarily stop interacting with the system [36]. It is often unclear whether such pauses reflect cognitive blockage or simply moments of reflective thinking. To address this limitation, future work could analyze the semantic patterns of users' recent ideas to assess whether their thinking is converging and becoming increasingly homogeneous. Incorporating this cognitive dimension into timing decisions may allow the system to deliver interventions that align more accurately with users' evolving ideation states.

Personalized assistance. Individual personality traits shape participants' sensitivity to shared information and influence their comfort with others' ideas, including their willingness to engage in critical evaluation. Future research could therefore explore more personalized information-sharing mechanisms that account for these individual differences, enabling the system to adapt both the frequency of idea sharing and the degree of content granularity to better support each user's ideation process.

Further discovery of ideation patterns. As an initial stage of this research track, *GraftMind* has preliminarily validated the effectiveness of this new group ideation setting. In future work, we aim to conduct large-scale comparative experiments in real usage scenarios. We seek to uncover users' behavior patterns within *GraftMind* and examine how they differ from those in traditional brainstorming settings. Such analysis will help us better understand the strengths and weaknesses of different ideation settings.

9 Conclusion

In this paper, we explored a novel group ideation setting that goes beyond existing group and individual brainstorming approaches. Drawing on the identified strengths of both solitary and collaborative ideation processes, we present *GraftMind*, a system that enables users to ideate in private workspaces while an AI mediator leverages collective ideas to provide real-time ideation assistance. Through our user study, we demonstrated the effectiveness of *GraftMind* in supporting group ideation. Our findings point to promising avenues for further investigation along this research track. We believe that as AI capabilities continue to advance, this mode of group ideation will offer options more attuned to users' habits and individual traits, fostering more efficient processes and higher-quality outcomes.

Acknowledgments

This research is partially funded by Research on HKUST(GZ) Practices (HKUST(GZ)-ROP2025002), the AI Research and Learning Base of Urban Culture under Project 2023WZJD008, the 2025 Guangdong Province Colleges Teaching Quality and Teaching Reform Project, the Guangdong Provincial Key Lab of Integrated Communication, the Project of DEGP (No. 2023KCXTD042), the Sensing and Computation for Ubiquitous Internet of Things (No. 2023B1212010007).

References

[1] Mohammed Almutairi, Charles Chiang, Yuxin Bai, and Diego Gomez-Zara. 2025. tAIfa: Enhancing Team Effectiveness and Cohesion with AI-Generated Automated Feedback. In *Proceedings of the 4th Annual Symposium on Human-Computer Interaction for Work (CHIWORK '25)*. Association for Computing Machinery, New York, NY, USA, Article 2, 25 pages. doi:10.1145/3729176.3729197

[2] Joshua Ashkinaze, Julia Mendelsohn, Li Qiwei, Ceren Budak, and Eric Gilbert. 2025. How AI Ideas Affect the Creativity, Diversity, and Evolution of Human Ideas: Evidence From a Large, Dynamic Experiment. In *Proceedings of the ACM Collective Intelligence Conference (CI '25)*. Association for Computing Machinery, New York, NY, USA, 198–213. doi:10.1145/3715928.3737481

[3] Ali Ateeq, Marwan Milhem, Mohammed Alzoraiki, Mohammed I. F. Dawwas, Shafeeq Ahmed Ali, and Ahmad Yahia Al Astal. 2024. The impact of AI as a mediator on effective communication: enhancing interaction in the digital age. *Frontiers in Human Dynamics* 6 (2024), 1467384. doi:10.3389/fhumd.2024.1467384

[4] Nicola Bartoli and Simone Benedetto. 2022. Driven by notifications – exploring the effects of badge notifications on user experience. *PLOS ONE* 17 (06 2022), 1–9. doi:10.1371/journal.pone.0270888

[5] Rhyse Bendell, Jessica Williams, Stephen M. Fiore, and Florian Jentsch. 2024. Individual and team profiling to support theory of mind in artificial social intelligence. *Scientific Reports* 14, 1 (02 June 2024), 12635. doi:10.1038/s41598-024-63122-8

[6] Midhad Blazevic, Lennart B. Sina, Cristian A. Secco, Melanie Siegel, and Kawa Nazemi. 2024. Real-Time Ideation Analyzer and Information Recommender. *Electronics* 13, 9 (2024), 1761. doi:10.3390/electronics13091761

[7] Luiz Fernando de Carvalho Botega and Jonny Carlos da Silva. 2020. An artificial intelligence approach to support knowledge management on the selection of creativity and innovation techniques. *Journal of Knowledge Management* 24, 5 (2020), 1107–1130. doi:10.1108/JKM-10-2019-0559

[8] Virginia Braun and Victoria Clarke. 2019. Reflecting on reflexive thematic analysis. *Qualitative Research in Sport, Exercise and Health* 11, 4 (2019), 589–597. doi:10.1080/2159676X.2019.1628806

[9] Luz Mabel Camacho. 1991. *Understanding Productivity Loss in Brainstorming Groups: A New Approach*. Ph. D. Dissertation. The University of Texas at Arlington, Arlington, TX, USA. https://www.proquest.com/openview/1cfe4691f34c6740aa3fb8a586a7ba12/1

[10] R. Nicholas Carleton, Donald R. McCreary, Peter J. Norton, and Gordon J. G. Asmundson. 2006. Brief fear of negative evaluation scale-revised. *Depression and Anxiety* 23, 5 (2006), 297–303. doi:10.1002/da.20142

[11] Jianlyu Chen, Shitao Xiao, Peitian Zhang, Kun Luo, Defu Lian, and Zheng Liu. 2024. M3-Embedding: Multi-Linguality, Multi-Functionality, Multi-Granularity Text Embeddings Through Self-Knowledge Distillation. In *Findings of the Association for Computational Linguistics: ACL 2024*, Lun-Wei Ku, Andre Martins, and Vivek Srikumar (Eds.). Association for Computational Linguistics, Bangkok, Thailand, 2318–2335. doi:10.18653/v1/2024.findings-acl.137

[12] Chih-Hsun Chuang, Susan E. Jackson, and Yuan Jiang. 2016. Can Knowledge-Intensive Teamwork Be Managed? Examining the Roles of HRM Systems, Leadership, and Tacit Knowledge. *Journal of Management* 42, 2 (2016), 524–554. doi:10.1177/0149206313478189

[13] Terry Connolly, Leonard M. Jessup, and Joseph S. Valacich. 1990. Effects of Anonymity and Evaluative Tone on Idea Generation in Computer-Mediated Groups. *Management Science* 36, 6 (June 1990), 689–703. doi:10.1287/mnsc.36.6.689

[14] Philippa Cordingley, Bart Crisp, and Paige Johns. 2019. *Future of Education and Skills 2030: Curriculum Analysis*. Technical Report EDU/EDPC(2019)16. OECD. https://one.oecd.org/document/EDU/EDPC(2019)16/en/pdf

[15] Sumit Kumar Dam, Choong Seon Hong, Yu Qiao, and Chaoning Zhang. 2024. A Complete Survey on LLM-based AI Chatbots. arXiv:2406.16937 [cs.CL] https://arxiv.org/abs/2406.16937

[16] Ernesto Damiani, Paolo Ceravolo, Fulvio Frati, Valerio Bellandi, Ronald Maier, Isabella Seeber, and Gabriela Waldhart. 2015. Applying recommender systems in collaboration environments. *Computers in Human Behavior* 51 (2015), 1124–1133. doi:10.1016/j.chb.2015.02.045

[17] Alan R Dennis and Joseph S Valacich. 1993. Computer brainstorms: More heads are better than one. *Journal of applied psychology* 78, 4 (1993), 531. doi:10.1037/0021-9010.78.4.531

[18] Alan R. Dennis and Michael L. Williams. 2003. Electronic Brainstorming: Theory, Research, and Future Directions. In *Group Creativity: Innovation through Collaboration*, Paul B. Paulus and Bernard A. Nijstad (Eds.). Oxford University Press, Oxford, UK, 160–178. doi:10.1093/acprof:oso/9780195147308.003.0008

[19] Kaiwalya Deshpande, Ayush Jain, Abhinav, Sushruta Mishra, Shalini Goel, and Rachit Garg. 2025. Empowering Real-Time Communication: A Seamless Chatting System Using Websocket. In *Innovative Computing and Communications (Lecture Notes in Networks and Systems, Vol. 1039)*, Aboul Ella Hassanien, Sameer Anand, Ajay Jaiswal, and Prabhat Kumar (Eds.). Springer, Singapore, 405–417. doi:10.1007/978-981-97-4152-6_29

[20] Michael Diehl and Wolfgang Stroebe. 1987. Productivity Loss in Brainstorming Groups: Toward the Solution of a Riddle. *Journal of Personality and Social Psychology* 53, 3 (1987), 497–509. doi:10.1037/0022-3514.53.3.497

[21] Jane L. E, Yu-Chun Grace Yen, Isabelle Yan Pan, Grace Lin, Mingyi Li, Hyoungwook Jin, Mengyi Chen, Haijun Xia, and Steven P. Dow. 2024. When to Give Feedback: Exploring Tradeoffs in the Timing of Design Feedback. In *Proceedings of the 16th Conference on Creativity & Cognition* (Chicago, IL, USA). Association for Computing Machinery, New York, NY, USA, 292–310. doi:10.1145/3635636.3656183

[22] Xiaocong Fan, John Yen, and Richard A. Volz. 2005. A theoretical framework on proactive information exchange in agent teamwork. *Artificial Intelligence* 169, 1 (2005), 23–97. doi:10.1016/j.artint.2005.06.005

[23] Xiaocong Fan, John Yen, Rui Wang, Shuang Sun, and Richard A. Volz. 2004. Context-Centric Proactive Information Delivery. In *Proceedings of the IEEE/WIC/ACM International Conference on Intelligent Agent Technology (IAT '04)*. IEEE Computer Society, USA, 31–37. doi:10.5555/1018422.1019837

[24] Mohammadreza Farrokhnia, Omid Noroozi, Yvette Baggen, Harm Biemans, and Armin Weinberger. 2025. Improving hybrid brainstorming outcomes with computer-supported scaffolds: Prompts and cognitive group awareness. *Computers & Education* 227 (2025), 105229. doi:10.1016/j.compedu.2024.105229

[25] Adrian Furnham and Tanya Yazdanpanahi. 1995. Personality differences and group versus individual brainstorming. *Personality and Individual Differences* 19, 1 (1995), 73–80. doi:10.1016/0191-8869(95)00009-U

[26] Amelia Wen Jiun Gan, Quoc Dang, Blaine Western, and Jose Castillo. 2023. AI-Mediated Group Ideation. In *Proceedings of the 41st International Conference on Education and Research in Computer Aided Architectural Design in Europe (eCAADe 2023, Vol. 2)*. eCAADe, Graz, Austria, 389–398. doi:10.52842/conf.ecaade.2023.2.389

[27] Joey F. George, George K. Easton, Jay F. Nunamaker, and Gregory B. Northcraft. 1990. A Study of Collaborative Group Work With and Without Computer-Based Support. *Information Systems Research* 1, 4 (December 1990), 394–415. doi:10.1287/isre.1.4.394

[28] Katusha Gerardini, Eleonora Diletta Sarcinella, Francesca Borghesi, Andrea Pozzi, Andrea Gaggioli, and Alice Chirico. 2025. Supporting Brainstorming with a Collaborative Platform or a Generative Artificial Intelligence Tool: An Exploratory Study. *Cyberpsychology, Behavior, and Social Networking* 28, 7 (2025), 524–533. doi:10.1089/cyber.2024.0518

[29] Erling Shane German. 2017. *An investigation of human-model interaction for model-centric decision-making*. Ph. D. Dissertation. Massachusetts Institute of Technology, Cambridge, MA, USA. http://hdl.handle.net/1721.1/111228

[30] Olga Goldenberg and Jennifer Wiley. 2011. Quality, Conformity, and Conflict: Questioning the Assumptions of Osborn's Brainstorming Technique. *The Journal of Problem Solving* 3, 2 (2011), 5. doi:10.7771/1932-6246.1093

[31] Olga Goldenberg and Jennifer Wiley. 2019. Individual and group brainstorming: Does the question matter? *Creativity Research Journal* 31, 3 (2019), 261–271. doi:10.1080/10400419.2019.1619399

[32] Gabriel Enrique Gonzalez, Dario Andres Silva Moran, Stephanie Houde, Jessica He, Steven I. Ross, Michael Muller, Siya Kunde, and Justin D. Weisz. 2024. Collaborative Canvas: A Tool for Exploring LLM Use in Group Ideation Tasks. In *Joint Proceedings of the ACM IUI 2024 Workshops (CEUR Workshop Proceedings, Vol. 3660)*. CEUR-WS.org, Greenville, South Carolina, USA, –. https://ceur-ws.org/Vol-3660/paper4.pdf

[33] Janice Griffiths-Hemans and Rajiv Grover. 2006. Setting the stage for creative new products: investigating the idea fruition process. *Journal of the Academy of Marketing Science* 34, 1 (December 2006), 27–39. doi:10.1177/0092070305281777

[34] Luca Grilli and Mattia Pedota. 2024. Creativity and artificial intelligence: A multilevel perspective. *Creativity and Innovation Management* 33, 2 (2024), 234–247. doi:10.1111/caim.12580

[35] J. P. Guilford. 1967. Creativity: Yesterday, Today, and Tomorrow. *Journal of Creative Behavior* 1, 1 (1967), 3–14. doi:10.1002/j.2162-6057.1967.tb00002.x

[36] Sophie Hall, Veerle M. Baaijen, and David Galbraith. 2024. Constructing theoretically informed measures of pause duration in experimentally manipulated writing. *Reading and Writing* 37, 2 (February 2024), 329–357. doi:10.1007/s11145-022-10284-4

[37] Jiaxuan Han and Ruqin Ren. 2025. Why unequal AI access enhances team productivity: the mediating role of interaction processes and cognitive diversity. *Frontiers in Psychology* 16 (September 2025), 1636906. doi:10.3389/fpsyg.2025.1636906

[38] Jessica He, Stephanie Houde, Gabriel E. Gonzalez, Darío Andrés Silva Moran, Steven I. Ross, Michael Muller, and Justin D. Weisz. 2024. AI and the Future of Collaborative Work: Group Ideation with an LLM in a Virtual Canvas. In *Proceedings of the 3rd Annual Meeting of the Symposium on Human-Computer Interaction for Work* (Newcastle upon Tyne, United Kingdom) *(CHIWORK '24)*. Association for Computing Machinery, New York, NY, USA, Article 9, 14 pages. doi:10.1145/3663384.3663398

[39] Xiaoxin He, Yijun Tian, Yifei Sun, Nitesh Chawla, Thomas Laurent, Yann LeCun, Xavier Bresson, and Bryan Hooi. 2024. G-retriever: Retrieval-augmented generation for textual graph understanding and question answering. *Advances in Neural Information Processing Systems* 37 (2024), 132876–132907. https:

//dl.acm.org/doi/10.5555/3737916.3742140

[40] Starr Roxanne Hiltz, Murray Turoff, and Kenneth Johnson. 1989. Experiments in group decision making, 3: disinhibition, deindividuation, and group process in pen name and real name computer conferences. *Decision Support Systems* 5, 2 (1989), 217–232. doi:10.1016/0167-9236(89)90008-0

[41] Patricia Hirsch, Luca Moretti, Sibel Askin, and Iring Koch. 2024. Examining the cognitive processes underlying resumption costs in task-interruption contexts: Decay or inhibition of suspended task goals? *Memory & Cognition* 52, 2 (2024), 271–284. doi:10.3758/s13421-023-01458-8

[42] Maura Kessel, Jan Kratzer, and Carsten Schultz. 2012. Psychological Safety, Knowledge Sharing, and Creative Performance in Healthcare Teams. *Creativity and Innovation Management* 21, 2 (2012), 147–157. doi:10.1111/j.1467-8691.2012.00635.x

[43] Abidullah Khan, Atefeh Shokrizadeh, and Jinghui Cheng. 2025. Beyond Automation: How Designers Perceive AI as a Creative Partner in the Divergent Thinking Stages of UI/UX Design. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 1105, 12 pages. doi:10.1145/3706598.3713500

[44] Runa Korde and Paul B Paulus. 2017. Alternating individual and group idea generation: Finding the elusive synergy. *Journal of Experimental Social Psychology* 70 (2017), 177–190. doi:10.1016/j.jesp.2016.11.002

[45] Nataliya Kosmyna, Eugene Hauptmann, Ye Tong Yuan, Jessica Situ, Xian-Hao Liao, Ashly Vivian Beresnitzky, Iris Braunstein, and Pattie Maes. 2025. Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task. arXiv:2506.08872 [cs.AI] https://arxiv.org/abs/2506.08872

[46] Werner H Kunz, Laszlo Sajtos, and Carlos Flavián. 2025. Beyond replacement: human-machine collaboration in the age of AI. *Journal of Service Management* 36, 4 (2025), 477–494. doi:10.1108/JOSM-04-2025-0194

[47] Jérémy La Scala, Natalia Bartłomiejczyk, Denis Gillet, and Adrian Holzer. 2025. Fostering innovation with generative AI: a study on human-AI collaborative ideation and user anonymity. https://infoscience.epfl.ch/entities/publication/fd9aa050-e4db-4541-98a2-07098f5153d5

[48] Helmut Lamm and Gisela Trommsdorff. 1973. Group versus individual performance on tasks requiring ideational proficiency (brainstorming): A review. *European journal of social psychology* 3, 4 (1973), 361–388. doi:10.1002/ejsp.2420030402

[49] Jr. Larson, James R. 2010. *In Search of Synergy in Small Group Performance* (1st ed.). Psychology Press, New York. doi:10.4324/9780203848784

[50] 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. doi:10.3390/make5040065

[51] Soohwan Lee, Seoyeong Hwang, and Kyungho Lee. 2024. Conversational Agents as Catalysts for Critical Thinking: Challenging Design Fixation in Group Design. arXiv:2406.11125 [cs.HC] https://arxiv.org/abs/2406.11125

[52] Sung-Chul Lee, Jaeyoon Song, Eun-Young Ko, Seongho Park, Jihee Kim, and Juho Kim. 2020. SolutionChat: Real-time Moderator Support for Chat-based Structured Discussion. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–12. doi:10.1145/3313831.3376609

[53] 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. https://dl.acm.org/doi/abs/10.5555/3495724.3496517

[54] Xiang Li, Lili Mou, Rui Yan, and Ming Zhang. 2016. StalemateBreaker: A Proactive Content-Introducing Approach to Automatic Human-Computer Conversation. arXiv:1604.04358 [cs.CL] https://arxiv.org/abs/1604.04358

[55] Reuben Luera, Ryan A. Rossi, Alexa Siu, Franck Dernoncourt, Tong Yu, Sungchul Kim, Ruiyi Zhang, Xiang Chen, Hanieh Salehy, Jian Zhao, Samyadeep Basu, Puneet Mathur, and Nedim Lipka. 2024. Survey of User Interface Design and Interaction Techniques in Generative AI Applications. arXiv:2410.22370 [cs.HC] https://arxiv.org/abs/2410.22370

[56] Yossi Maaravi, Ben Heller, Yael Shoham, Shay Mohar, and Baruch Deutsch. 2021. Ideation in the digital age: literature review and integrative model for electronic brainstorming. *Review of Managerial Science* 15, 6 (2021), 1431–1464. doi:10.1007/s11846-020-00400-5

[57] Susanne Melker, Emma Gabrils, Victor Villavicencio, Montathar Faraon, and Kari Rönkkö. 2025. Artificial intelligence for design education: a conceptual approach to enhance students' divergent and convergent thinking in ideation processes. *International Journal of Technology and Design Education* 35 (November 2025), 1871–1899. doi:10.1007/s10798-025-09964-3

[58] Melkamu Abay Mersha and Jugal Kalita. 2025. Semantic-Driven Topic Modeling for Analyzing Creativity in Virtual Brainstorming. arXiv:2509.16835 [cs.CL] https://arxiv.org/abs/2509.16835

[59] Jessica R Mesmer-Magnus and Leslie A DeChurch. 2009. Information sharing and team performance: a meta-analysis. *Journal of Applied Psychology* 94, 2 (2009), 535. doi:10.1037/a0013773

[60] Alberto Monge Roffarello, Luigi De Russis, and Massimiliano Pellegrino. 2024. Digital Wellbeing Lens: Design Interfaces That Respect User Attention. In *Proceedings of the 2024 International Conference on Advanced Visual Interfaces* (Arenzano, Genoa, Italy) *(AVI '24)*. Association for Computing Machinery, New York, NY, USA, Article 51, 5 pages. doi:10.1145/3656650.3656674

[61] Michael Muller, Stephanie Houde, Gabriel Gonzalez, Kristina Brimijoin, Steven I. Ross, Dario Andres Silva Moran, and Justin D. Weisz. 2024. Group Brainstorming with an AI Agent: Creating and Selecting Ideas. In *Proceedings of the International Conference on Computational Creativity*. Association for Computational Creativity, Jönköping, Sweden, 10. https://computationalcreativity.net/iccc24/papers/ICCC24_paper_18.pdf

[62] Bernard Nijstad and Wolfgang Stroebe. 2006. How the Group Affects the Mind: A Cognitive Model of Idea Generation in Groups. *Personality and social psychology review : an official journal of the Society for Personality and Social Psychology, Inc* 10 (02 2006), 186–213. doi:10.1207/s15327957pspr1003_1

[63] Bernard A Nijstad, Carsten KW De Dreu, Eric F Rietzschel, and Matthijs Baas. 2010. The dual pathway to creativity model: Creative ideation as a function of flexibility and persistence. *European review of social psychology* 21, 1 (2010), 34–77. doi:10.1080/10463281003765323

[64] Moeka Nomura, Takayuki Ito, and Shiyao Ding. 2024. Towards Collaborative Brain-storming among Humans and AI Agents: An Implementation of the IBIS-based Brainstorming Support System with Multiple AI Agents. In *Proceedings of the ACM Collective Intelligence Conference* (Boston, MA, USA) *(CI '24)*. Association for Computing Machinery, New York, NY, USA, 1–9. doi:10.1145/3643562.3672609

[65] Alex F. Osborn. 1953. *Applied Imagination: Principles and Procedures of Creative Problem Solving*. Charles Scribner's Sons, New York, NY.

[66] Srishti Palani, David Ledo, George Fitzmaurice, and Fraser Anderson. 2022. "I don't want to feel like I'm working in a 1960s factory": The Practitioner Perspective on Creativity Support Tool Adoption. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 379, 18 pages. doi:10.1145/3491102.3501933

[67] Paul B Paulus, Jonali Baruah, and Jared B Kenworthy. 2018. Enhancing collaborative ideation in organizations. *Frontiers in psychology* 9 (2018), 2024. doi:10.3389/fpsyg.2018.02024

[68] Paul B Paulus, Daniel S Levine, Vincent Brown, Ali A Minai, and Simona Doboli. 2010. Modeling ideational creativity in groups: Connecting cognitive, neural, and computational approaches. *Small Group Research* 41, 6 (2010), 688–724. doi:10.1177/1046496410369561

[69] Alain Pinsonneault and Nelson Heppel. 1997. Anonymity in group support systems research: A new conceptualization, measure, and contingency framework. *Journal of Management Information Systems* 14, 3 (1997), 89–108. http://www.jstor.org/stable/40398279

[70] Tom Postmes and Martin Lea. 2000. Social processes and group decision making: anonymity in group decision support systems. *Ergonomics* 43, 8 (2000), 1252–1274. doi:10.1080/00140130050084978

[71] Leonard Przybilla, Luka Baar, Manuel Wiesche, and Helmut Krcmar. 2019. Machines as Teammates in Creative Teams: Digital Facilitation of the Dual Pathways to Creativity. In *Proceedings of the 2019 on Computers and People Research Conference* (Nashville, TN, USA) *(SIGMIS-CPR '19)*. Association for Computing Machinery, New York, NY, USA, 94–102. doi:10.1145/3322385.3322402

[72] Kevin Pu, K. J. Kevin Feng, Tovi Grossman, Tom Hope, Bhavana Dalvi Mishra, Matt Latzke, Jonathan Bragg, Joseph Chee Chang, and Pao Siangliulue. 2025. IdeaSynth: Iterative Research Idea Development Through Evolving and Composing Idea Facets with Literature-Grounded Feedback. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 145, 31 pages. doi:10.1145/3706598.3714057

[73] Kevin Pu, Daniel Lazaro, Ian Arawjo, Haijun Xia, Ziang Xiao, Tovi Grossman, and Yan Chen. 2025. Assistance or Disruption? Exploring and Evaluating the Design and Trade-offs of Proactive AI Programming Support. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 152, 21 pages. doi:10.1145/3706598.3713357

[74] Gerard J Puccio, Mary C Murdock, and Marie Mance. 2005. Current developments in creative problem solving for organizations: A focus on thinking skills and styles. *Korean Journal of Thinking and Problem Solving* 15, 2 (2005), 43. https://psycnet.apa.org/record/2005-13655-004

[75] Vicky L Putman and Paul B Paulus. 2009. Brainstorming, brainstorming rules and decision making. *The Journal of creative behavior* 43, 1 (2009), 29–40. doi:10.1002/j.2162-6057.2009.tb01304.x

[76] Peinuan Qin, Chi-Lan Yang, Jingshu Li, Jing Wen, and Yi-Chieh Lee. 2025. Timing Matters: How Using LLMs at Different Timings Influences Writers' Perceptions and Ideation Outcomes in AI-Assisted Ideation. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 25, 16 pages. doi:10.1145/3706598.3713146

[77] Alireza Rezazadeh, Zichao Li, Ange Lou, Yuying Zhao, Wei Wei, and Yujia Bao. 2025. Collaborative Memory: Multi-User Memory Sharing in LLM Agents with Dynamic Access Control. arXiv:2505.18279 [cs.MA] https://arxiv.org/abs/2505.18279

[78] John R Rossiter and Gary L Lilien. 1994. New "brainstorming" principles. *Australian Journal of Management* 19, 1 (1994), 61–72. doi:10.1177/031289629401900104

[79] Christin Schulze, Ada Aka, Daniel M. Bartels, Stefan F. Bucher, Jake R. Embrey, Todd M. Gureckis, Gerald Häubl, Mark K. Ho, Ian Krajbich, Alexander K. Moore, Gabriele Oettingen, Jonathan D. K. Ongchoco, Ryan Oprea, Nicholas Reinholtz, and Ben R. Newell. 2025. A timeline of cognitive costs in decision-making. *Trends in Cognitive Sciences* 29, 9 (2025), 827–839. doi:10.1016/j.tics.2025.04.004

[80] Orit Shaer, Angelora Cooper, Andrew L. Kun, and Osnat Mokryn. 2024. Toward Enhancing Ideation through Collaborative Group-AI Brainwriting. In *Joint Proceedings of the ACM IUI 2024 Workshops (CEUR Workshop Proceedings, Vol. 3660)*. CEUR-WS.org, Greenville, South Carolina, USA, 13 pages. https://ceur-ws.org/Vol-3660/paper13.pdf

[81] Orit Shaer, Angelora Cooper, Osnat Mokryn, Andrew L Kun, and Hagit Ben Shoshan. 2024. AI-Augmented Brainwriting: Investigating the use of LLMs in group ideation. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 1050, 17 pages. doi:10.1145/3613904.3642414

[82] Patrick Chien-Pai Shih. 2011. *Brainstorming beyond the laboratory: idea generation practices in software development firms.* Ph. D. Dissertation. University of California, Irvine, USA. Advisor(s) Olson, Gary M. and Redmiles, David F.

[83] Joon Gi Shin, Janin Koch, Andrés Lucero, Peter Dalsgaard, and Wendy E. Mackay. 2023. Integrating AI in Human-Human Collaborative Ideation. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) *(CHI EA '23)*. Association for Computing Machinery, New York, NY, USA, Article 355, 5 pages. doi:10.1145/3544549.3573802

[84] Dean Keith Simonton. 1999. *Origins of Genius: Darwinian Perspectives on Creativity.* Oxford University Press, New York, NY. doi:10.1093/oso/9780195128796.001.0001

[85] Diane H. Sonnenwald and Leah A. Lievrouw. 1997. Collaboration during the design process: a case study of communication, information behavior, and project performance. In *Proceedings of an International Conference on Information Seeking in Context* (Tampere, Finland) *(ISIC '96)*. Taylor Graham Publishing, GBR, 179–204. doi:10.5555/267190.267201

[86] Wolfgang Stroebe and Michael Diehl. 1994. Why groups are less effective than their members: On productivity losses in idea-generating groups. *European review of social psychology* 5, 1 (1994), 271–303. doi:10.1080/14792779543000084

[87] Oguzhan Topsakal and Tahir Cetin Akinci. 2023. Creating Large Language Model Applications Utilizing LangChain: A Primer on Developing LLM Apps Fast. In *Proceedings of the 5th International Conference on Applied Engineering and Natural Sciences (ICAENS)*. All Sciences Proceedings, Konya, Turkey, 1050–1056. doi:10.59287/icaens.1127

[88] Ryan Watkins and Soheil Human. 2023. Needs-aware artificial intelligence: AI that 'serves [human] needs'. *AI and Ethics* 3, 1 (2023), 49–52. doi:10.1007/s43681-022-00181-5

[89] Anita Williams Woolley. 2025. Generative AI and collaboration: opportunities for cultivating collective intelligence. *Journal of Organization Design* in press (2025), 6 pages. doi:10.1007/s41469-025-00199-z Published online: 10 October 2025.

[90] An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu. 2025. Qwen3 Technical Report. arXiv:2505.09388 [cs.CL] https://arxiv.org/abs/2505.09388

[91] John Zerilli, Umang Bhatt, and Adrian Weller. 2022. How transparency modulates trust in artificial intelligence. *Patterns* 3, 4 (2022), 100455. doi:10.1016/j.patter.2022.100455

[92] Yanzhao Zhang, Mingxin Li, Dingkun Long, Xin Zhang, Huan Lin, Baosong Yang, Pengjun Xie, An Yang, Dayiheng Liu, Junyang Lin, Fei Huang, and Jingren Zhou. 2025. Qwen3 Embedding: Advancing Text Embedding and Reranking Through Foundation Models. arXiv:2506.05176 [cs.CL] https://arxiv.org/abs/2506.05176

[93] Zheng Zhang, Weirui Peng, Xinyue Chen, Luke Cao, and Toby Jia-Jun Li. 2025. 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 (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 147, 22 pages. doi:10.1145/3706598.3713289

[94] Michelle Zhao, Fade R Eadeh, Thuy-Ngoc Nguyen, Pranav Gupta, Henny Admoni, Cleotilde Gonzalez, and Anita Williams Woolley. 2023. Teaching agents to understand teamwork: Evaluating and predicting collective intelligence as a latent variable via Hidden Markov Models. *Computers in Human Behavior* 139 (2023), 107524. doi:10.1016/j.chb.2022.107524

[95] Xiang Zhou, Hong-Kun Zhai, Bibi Delidabieke, Hui Zeng, Yu-Xin Cui, and Xue Cao. 2019. Exposure to ideas, evaluation apprehension, and incubation intervals in collaborative idea generation. *Frontiers in psychology* 10 (2019), 1459. doi:10.3389/fpsyg.2019.01459